MODERN VHF/UHF FRONT-END DESIGN

Ian White, G3SEK*

Ian White is a self-confessed vhf/uhf dxer. There is no known cure for this condition, a side-effect of which is an interest in any related technical aspect of amateur radio. When the bands are closed and the moon is not up, he enjoys writing (hence this article), being secretary of the Vale of White Horse ARS, and being musician for the Abingdon Traditional Morris Dancers.

INTRODUCTION

VHF/UHF DX OPERATING is unique to amateur radio. Unlike any other users of those frequencies, we communicate over long distances using ssb and cw, with continuous tuning and not fixed channels. In other words, we treat the low-frequency ends of our vhf, uhf and even microwave bands just like hf. So it is hardly surprising that our standards for vhf, uhf and microwave dx receivers are strongly based on the kind of performance we expect from hf dx receivers. In particular we want to be able to copy weak dx signals, even if there are plenty of strong local signals on the band at the same time. The extra dimension on our bands above 50MHz is that the background noise is much lower. An hf signal of a few microvolts is often down in the noise, but at vhf and above we can communicate with signal levels down to a few *nano*volts—if the other strong signals will let us, that is!

What makes a receiver good (or bad) at handling both weak and strong signals is its "front-end", the stages between the antenna and the selective i. f. filter. Your receiver's front-end has to amplify and frequency-change the weak wanted signal without introducing too much extra noise, yet it also has to cope with an enormous number of unwanted signals. Fortunately, you never hear most of these other signals, because they are eliminated once they reach the selective i.f. filter. But very strong signals may already have overloaded the front-end before they get that far. And front-end overload can seriously interfere with the weak dx signals that you most want to hear.

A certain amount of front-end gain is essential in a sensitive receiver. Yet too much gain will make strong signals build up to levels that overload the front-end. To achieve the optimum *overall* performance, you have to make a very careful trade-off between sensitivity and strong-signal handling.

This article is about front-end *design*. Starting from basics I will explain what we amateurs need from modern vhf/uhf front-ends, and how to go about getting it. We certainly *won't* get what we need by patching together a few off-the-peg circuits featuring this year's fashionable devices. In the 'eighties that approach is outdated and—in the worst sense—amateurish. These days we can use the power of scientific calculators and home computers to predict the performance of our designs *before* we build them. Having shown you how the calculations work, I will describe a program that lets you experiment with different designs until you have just what you need. And then, knowing exactly how your new front-end should perform, you can go ahead confidently and build it.

Since vhf/uhf dxing is unique to amateur radio, the people at the leading edge in front-end design are all radio amateurs, even if some of them (not me) also do it for a living! I hope this article will bring a few more people up towards that leading edge. More than that, I want the article to tell *every* vhf/uhf dxer what to expect from a modern front-end. And if the developments are taken up by the big transceiver manufacturers, then maybe one day we can *all* enjoy the benefits of modern vhf/uhf front-end design.

*52 Abingdon Road, Drayton, Abingdon, Oxon OX14 4HP

PART 1. WEAK-SIGNAL PERFORMANCE

THERE ARE many ways to specify the weak-signal performance of a receiver, but underlying them all is the basic concept of noise and temperature. Although you may have heard the term before, you may not realize how useful noise temperature can be in drawing together many different aspects of receiver design. This part of the article will explain what noise temperature means, and relate it to some other ways of specifying receiver sensitivity. Armed with these basic ideas, we can decide on some reasonable targets for receiver sensitivity on the various vhf/uhf bands. Then we can see if the front-ends of commercial transceivers measure up to our requirements (no prizes for guessing the answer!). Finally I will describe how to analyse system noise performance.

Noise temperature

Any electrical conductor contains electrons which are free to move around. At normal temperatures, electrons in a conductor are in random motion, rushing back and forth but, on average, getting nowhere. This random thermal movement of electrons constitutes a fluctuating current, which can be detected as random noise. Converted to audio frequencies, random noise is the rushing sound your receiver makes when tuned to an unoccupied frequency. (You may also hear man-made noise which is obviously not random; eg, ignition noise with its regular repetition rate.) All random thermal movement ceases at a temperature of absolute zero (−273·16°C), so no thermal noise is produced. At any temperature above absolute zero, the thermal noise power generated in a conductor is proportional to its physical temperature, measured on the "absolute" temperature scale in units of kelvins (K). Thermal noise is spread evenly over the entire electromagnetic spectrum, so the noise power detected by a receiver is proportional to the spectral bandwidth from which the noise is "collected". The basic formula relating noise power P_N (Watts) to noise temperature T (K) and bandwidth B (Hz) is simply:

$$P_N = kTB \quad \ldots\ldots (1)$$

where k is called Boltzmann's constant and has the value $1{\cdot}38 \times 10^{-23}$ J/K (joules/kelvin). For example, the thermal noise power generated in a resistor at a temperature of 290K (around room temperature) is:

$$1{\cdot}38 \times 10^{-23} \times 290 \times B \text{ watts;}$$

ie, $4{\cdot}00 \times 10^{-21}$W in every hertz of bandwidth B, across the entire spectrum. Two things to note: the noise power is independent of the ohmic value of the resistor, and the value calculated in equation (1) is the power that would be delivered into a matched load having the same impedance as the resistor.

Another example: if a 50Ω resistor at 290K was connected to the input terminals of a noise-free receiver, with a 50Ω input impedance and a bandwidth of 2·5kHz (Fig 1), what would be the noise voltage? The noise power from (1) is $1{\cdot}38 \times 10^{-23} \times 290 \times 2{,}500 = 1{\cdot}00 \times 10^{-17}$W. The voltage across 50Ω is obtained simply from $P = V^2/R$, and is $(1{\cdot}035 \times 10^{-17} \times 50) = 22{\cdot}7$nV.

Resistors are not the only electronic devices which generate noise [1]. Imagine an rf preamplifier (Fig 2) connected between the 50Ω resistor and the noise-free receiver. The noise power delivered to the receiver now has

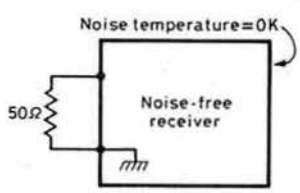

Fig 1. Even a 50Ω resistor generates some noise

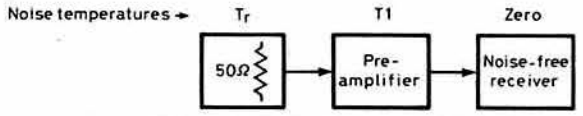

Fig 2. The effect of adding a preamplifier

two components: one is the amplified thermal noise from the resistive noise source, and the other is noise generated in the preamplifier itself. The noise contribution from the preamplifier can be expressed simply as a power in watts; or more subtly in terms of an imaginary increase in the temperature of the source resistor that would be needed to produce the same noise power. This imaginary temperature increase is known as the equivalent noise temperature of the preamplifier (T1 in Fig 2). A perfect device would have a noise temperature of zero K so it would contribute no extra noise. All real devices have a positive noise temperature, and the higher their equivalent noise temperature, the more noise they contribute.

The practical advantage from thinking in terms of equivalent noise temperature is that we have a common basis for measuring random noise arising from just about anything from a GaAsfet to a galaxy. The noise does not even have to be truly thermal in origin. The equivalent noise temperature of an electronic device is rarely the same as its physical temperature (unless the device is a resistor). For example, a modern mosfet amplifier at room temperature (around 290K) may have a noise temperature as low as 50K at 144MHz, while a very noisy device can have a noise temperature well above 1,000K without glowing white-hot! Noise from an active device like a transistor (or a galaxy) is also frequency-dependent, unlike true thermal noise, so although the noise temperature may be taken as constant within an amateur band, it will probably be different on another band.

Compared with other ways of measuring noise, noise temperature has some very useful properties. It is independent of bandwidth (unlike noise power, which is proportional to bandwidth), and it is independent of the gain or loss of the device generating the noise. Most useful of all, the noise temperatures of different components in a total system can be added directly to give a system noise temperature.

input. The excess noise temperature is simply T2/G1, where G1 is the power gain of the first stage, so the system noise temperature T_{SYS} is given by:

$$T_{SYS} = t_R + T1 + T2/G1 \quad \ldots\ldots\ldots\ldots\ldots\ldots (2)$$

Dividing T2 by G1 is called "referring T2 to the input". And we can continue in this way, adding on extra stages (Fig 4a) and referring their noise contributions back to the input, so that:

$$T_{SYS} = T_R + T1 + T2/G1 + T3/(G1 \times G2) + T4/(G1 \times G2 \times G3) \ldots (3)$$

The noise temperature T3 of the third stage is referred back to the input by dividing it by the cumulative gain up to that point (G1 × G2 of the two previous stages), and so on. Notice how simply the noise temperatures add up. Notice also that the power gains G1, G2 etc are pure ratios (eg, "a power gain of 15 times") and are not in decibels.

Let's go back to the simplest two-stage system (Fig 3) and look at equation (2) in more detail. There are limits to how low the system noise temperature can be. For a start, it cannot be less than T_R, the noise temperature of the source from which the system gets its signals. The noise temperature T1 of the first stage also adds directly to the system total. The situation is different for the later stages, though. T_{SYS} can be brought close to its lower limit of (T_R + T1) if T2 is made small, or G1 large, since either will tend to make the last term (T2/G1) smaller. The same applies to the multi-stage system in Fig 4(a) and equation (3), for the cumulative effect of G1, G2 etc on the bottom lines of the fractions can quickly make the contributions of later stages negligible. To show this effect in practice, let's put some typical performance figures into the system of Fig 4(a), to give Fig 4(b). In a real multistage system we can say that the last stage is "the rest of the system", so we no longer need that fictional noise-free receiver in Figs 1-4(a). Then using equation (3) we have:

$$T_{SYS} = 300K + 170K + 250K/10 + 500K/(10\times 10) + 1{,}000K/(10\times 10\times 10)$$
$$= 300K + 170K + 25K + 5K + 1K$$
$$= 501K$$

The cumulative gain makes the noise contributions of the later stages very small, even though the individual noise temperatures of those stages are quite high. Mixers behave in exactly the same way as far as these equations are concerned, and even lossy mixers such as diode rings can be treated in the same way (see later).

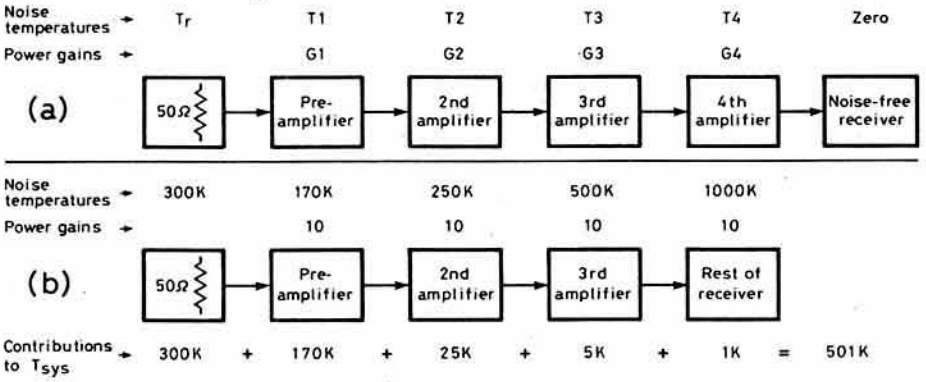

Fig 4. Noise analysis of more stages: (a) in symbols; (b) with some typical numbers

System noise temperature

We have seen that the noise temperature of a simple rf preamplifier fed from a 50Ω resistor (Fig 2), is the sum of the noise temperature of the resistor T_R (equal to its physical temperature because it is a resistor) plus the equivalent noise temperature T1 of the preamplifier. This sum is a simple form of system noise temperature. Now let us add to the system another stage whose noise temperature is T2 (Fig 3). We can calculate the noise temperature of this new system by finding out what excess noise the second stage contributes, in terms of excess equivalent noise temperature at the

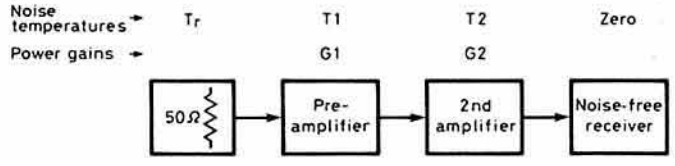

Fig 3. A simple noise analysis

We also need to know how to deal with such loss-introducing devices as cables, attenuators and filters. These degrade the system noise temperature, mainly by attenuating the wanted signal. There is also another less obvious source of degradation, due to the additional thermal noise generated in the resistive part of the losses. This second effect is little known and often neglected, but it can be important in the ultra-low-noise systems we are dealing with here. The complete formula for the effect of lossy components on system noise temperature is:

$$T = T_F L + (L - 1)T_P \quad \ldots\ldots\ldots\ldots\ldots\ldots (4)$$

where:

T is the noise temperature at the input of the lossy component;

L is the loss expressed as a power ratio (ie, L = 1 for a loss-less system, going up to infinity for infinite loss);

T_F is the cumulative noise temperature of all following stages, referred to the output of the lossy device; and

T_P is the physical temperature of the lossy device.

The first term in equation (4) is the attenuation term, and the second is the term for extra thermal noise, which depends on the physical temperature of the lossy device. Note that we have been assuming that the losses are resistive and occur in a matched system. If the losses are due to a reactive mismatch, only the first term in equation (4) is relevant.

Noise factor and noise figure

Noise temperature is the fundamental concept underlying every aspect of noise in receiving systems. So how are the more familiar-sounding terms "noise factor" and "noise figure" related to noise temperature?

Noise factor F is related to noise temperature T by a standard definition [2]:

$$F = 1 + T/290$$

290K is close to room temperature, so noise factor is a measure of how the device or system noise compares with thermal noise under typical conditions of use. From this definition the perfect system has a noise factor of 1, and all real noise factors are greater than 1. Even more familiar than noise factor is noise figure N, which is simply noise factor expressed in decibels.

$$N = 10 \log_{10}(1 + T/290) \quad \ldots\ldots\ldots\ldots\ldots\ldots (5)$$

Since the minimum possible noise factor is 1, the minimum possible noise figure is 0dB. Unfortunately, noise factor and noise figure are easily confused, and there is little need to become familiar with both. In this series I will be using either noise temperature (and always with its units of kelvins) or noise figure (and always with its units of decibels).

Noise figure (nf) and noise temperature each have their advantages. The attenuation of a lossy component (in decibels) can be added to the nf of the following amplifier or mixer stage (also in decibels) to give a new, larger nf for the two stages combined [3]. But nfs of cascaded amplifier or mixer stages do *not* add up in the same simple way as noise temperatures do. And there is another more serious problem about nf: it can only deal with noise in receivers, and receivers are not the whole story.

Antenna noise temperatures

The receiver is only part of the receiving system as a whole. An equally important contributor to the complete system noise temperature T_{SYS} is the antenna, which can be brought into the common framework by giving it an equivalent noise temperature.

The noise temperature of an antenna is determined by the rf environment that it can "see", and not by its physical temperature. For example, a 144MHz beam pointed at the sky for satellite work may be picking up only a very small amount of galactic noise, which can be represented by an equivalent noise temperature as low as 150–200K [4]. On the other hand, if the antenna is picking up thermal noise from its earthly surroundings at physical temperatures of 270–300K or whatever, that will also be its noise temperature. The usual situation is somewhere in between: at 144MHz flat ground is a fairly good reflector of radiation and a correspondingly poor emitter; ie, the ground acts mostly like a mirror. So a horizontal antenna sees either "cold" sky above the horizon or the sky image reflected by the ground, and the warm ground itself makes only a small contribution to the antenna's equivalent noise temperature. Pick-up of man-made noise may be much more important than thermal or galactic noise, but, so long as the man-made noise is virtually random, this too can be expressed as an equivalent noise temperature. (See what I meant about the advantages of thinking about lots of different things in the same equivalent terms?) The noise temperature of a 144MHz antenna can thus vary between less than 250K in satellite work to 1,000K or more in an urban environment full of man-made noise.

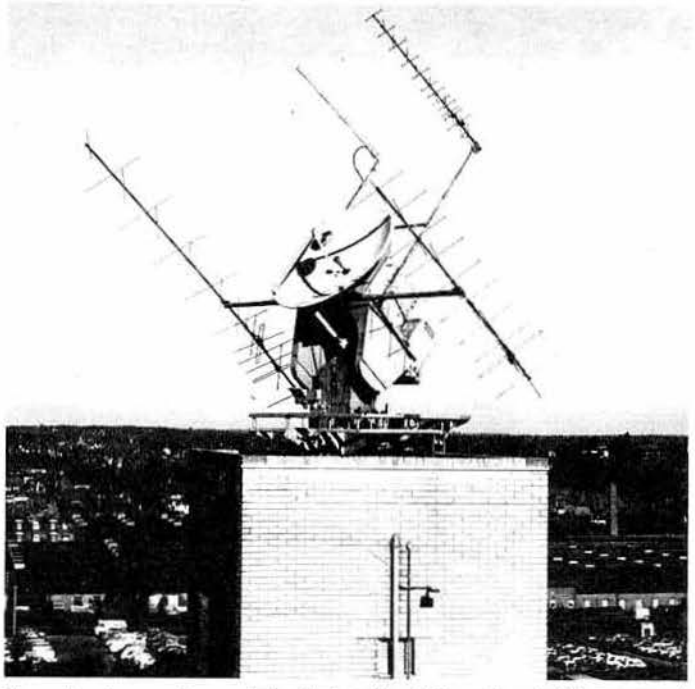

The noise temperatures of the University of Surrey's satellite command station antennas would be much higher if they were aimed at the horizon across the city of Guildford

For terrestrial working on 144MHz, 200K is a reasonable minimum estimate for the antenna noise temperature. At higher frequencies, galactic and man-made noise tend to decrease, but the ground and nearby objects become poorer reflectors and better emitters of thermal noise. At 1·3GHz and above, an antenna pointed at the ground will have a noise temperature corresponding closely to the physical ground temperature, while a normal horizontal antenna will see approximately half warm ground and half cold sky, so its noise temperature will be about 150K. The same noise temperature can also be assumed for 432MHz antennas aimed at the horizon.

Usable sensitivity

Knowing that system noise temperature can be divided into two contributions, from the antenna and from the receiver, we can now tackle the question: how much sensitivity do we *really* need at vhf/uhf? In olden days, when rf devices were noisy and the vhf/uhf bands were empty, the answer was simple: we always needed more sensitivity, simply to improve the chances of ever hearing a signal at all! Times change; and in more ways than one. The low-noise rf devices that we dreamed of a decade ago are here, but so too are crowded bands full of strong signals. On today's vhf/uhf bands, receiver overload due to too much front-end gain can be a worse problem than lack of sensitivity.

Modern front-end design calls for a deliberate trade-off between sensitivity and strong-signal handling, to achieve the optimum *overall* balance of performance. In an optimized front-end, system noise temperature becomes something to be decided upon and designed for, instead of naively being "as low as possible". But what noise temperature should we choose?

First of all, remember that the system noise temperature cannot possibly be lower than the antenna noise temperature, even with a perfectly noiseless receiver. Noise from any practical receiver must degrade the signal-to-noise ratio of a weak signal to some extent, and the choice of receiver noise temperature can be framed in terms of how much degradation is allowable.

Let's try a few numbers. Consider a 144MHz receiver connected to an antenna with a noise temperature of 200K, the "typical" figure suggested earlier. If the receiver noise temperature is also 200K (nf = 2·28dB), the system noise temperature is 400K, and the noise power is doubled. Thus the signal : noise ratio of an incoming signal would be degraded by 3dB, compared with the unachievable noiseless receiver. To reduce the degradation to 2dB would mean reducing the system noise temperature to 317·7K. The antenna noise temperature is 200K and cannot be changed, so the receiver noise temperature must come down to 117·7K (nf = 1·48dB). To reduce the signal : noise degradation to only 1dB would call for a receiver noise temperature of 52·4K (nf = 0·72dB); such low noise temperatures can of course be achieved, but only by sacrificing other aspects of receiver performance, particularly strong-signal handling.

From these trial calculations, it seems that if we begin with a receiver noise temperature which is equal to the antenna noise temperature, the real scope for further improving the signal : noise ratio is only a decibel or two. Such a small difference may scarcely be noticeable on the air. And any such improvement may involve unacceptable sacrifices in the receiver's ability to handle strong signals. In other words, *if the receiver noise temperature is as low as the noise temperature of the antenna to which it is connected, the receiver is as sensitive as it really needs to be.*

Referring back to the typical antenna noise temperatures, the design-target receiver noise temperatures (nfs in brackets) for ordinary terrestrial dx-chasing would thus be 200K (say, 2·2dB) at 144MHz, and 150K (1·8dB) at 432MHz and above. These are not hard-and-fast rules, merely recommendations in round figures, given the current state of the art of vhf/uhf front-end design. The optimum balance between sensitivity and strong-signal handling may change as techniques improve, though the lower limit of system noise temperature will not: that limit is set by the antenna and its environment, and not by the receiver.

Sensitivity of commercial equipment

How does commercial equipment measure up to these standards of sensitivity? Not particularly well. Commercial transceivers often seem rather deaf, and their sensitivity to weak signals can be noticeably improved by a preamp or a replacement front-end.

In order to compare the published sensitivity specifications of

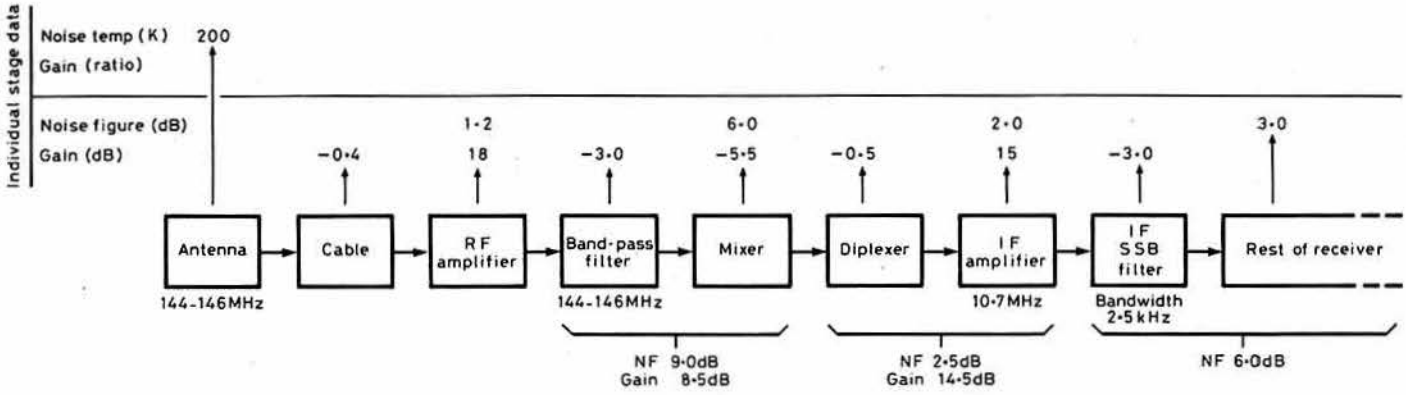

Fig 5. Block diagram of 144MHz front-end, showing input data for individual stages, and how some pairs of stages can be combined for analysis

commercial transceivers with the design targets recommended above, we first need to understand the way those specifications are worded. The most common way of expressing receiver sensitivity at hf is in terms of "$x\mu$V for a signal : noise ratio of ydB". Regrettably, the big transceiver manufacturers specify the sensitivity of their vhf and uhf equipment in that way too: regrettably, because out of all the possible ways of doing it, that one is probably the worst!

One criticism of the "microvolt" method is that it offers several possible grounds for misunderstandings, for getting the measurement wrong, or for presenting the figures so that they look good. For example, the signal level can be expressed either as a potential difference (pd) in a 50Ω system or as an open-circuit emf; there is a 6dB difference. "Signal : noise ratio" may really mean "(signal + noise) : noise ratio", which may differ by a decibel or so. Also, receiver noise is proportional to the bandwidth, so the sensitivity defined by the microvolt method will be better when using a narrow cw bandwidth than when the ssb filter is selected. It is extremely difficult to make measurements on ultra-sensitive vhf and uhf front-ends using traditional signal-generator methods. The possibilities for mistakes are enormous, and usually lead to over-optimistic results.

But the most serious criticism of the microvolt specifications for vhf/uhf equipment is that they are often almost worthless! For example, the manual for my 144MHz transceiver says the sensitivity should be "better than 0·5μV for 10dB s : n". This specification can be translated into noise temperature or nf by the method given in note [5]. Making the most favourable assumptions, the translated specification would read "noise temperature better than 3,000K" or "better than 11dB nf". Having just read that the design target for a 144MHz front-end should be a noise temperature around 200K (2·2dB nf), you can see how little the manufacturer is really promising.

While the sensitivity of the average commercial transceiver is often much better than the specification implies, most vhf and uhf transceivers *are* distinctly hard-of-hearing—hence the brisk trade in preamps, or even complete replacement front-ends! We still buy the transceivers for their other good points, but until consumers are given meaningful sensitivity specifications, the manufacturers have little incentive to get to grips with the problems of front-end design.

Final examples

Having, I hope, provoked a few thoughts about what receiver sensitivity really means, I will round off this part of the article with a complete example of noise-temperature analysis. The system we will consider is a 144MHz single-conversion front-end, all the way back from the i. f. filter to the antenna. In a later article I will return to this example to analyse its strong-signal performance too.

We start by drawing a block diagram of all the separate stages of the front-end, with the signal flowing from left to right (Fig 5). Don't worry about what kind of circuitry goes inside the blocks; what we are concerned with at present is simply the performance of each block in terms of noise temperature and power gain. The performance of the system as a whole will be calculated from the data on the individual stages, so we begin by writing what we know about each stage in the space above it. The data come from such sources as manufacturers' data sheets, magazine articles, or measurements on prototypes, and are generally given in decibel units of noise figure and gain. These will have to be converted to noise temperature and gain ratio.

We can simplify the calculations before we even start. In Fig 5 there are four places where an amplifier or mixer has a lossy stage ahead of it, and in each case the two could be combined [3] to save effort. The "rest of receiver" has an nf of 3dB, and ahead of it is the i. f. filter with a loss of 3dB, so these two can be combined into a single stage with an nf of 6dB. The first i. f. amplifier can be combined with the diplexer, and the mixer with the bandpass filter. The rf amplifier *could* be combined with the cable, though it is better to keep these two stages separate.

The simplified block diagram is in the middle of Fig 6, with the individual-stage data above. Below the block diagram will go the results for the system. If you want to follow through this worked example, draw up a similar worksheet for yourself. The input data from Fig 5 were mostly in decibel form, so check that you can do the conversions into noise temperatures and gain ratios. The relevant formulas are:

$$T = 290\ (\text{antilog}_{10}(N/10) - 1) \quad \ldots\ldots\ldots \text{(from equation 5)}$$

for converting nf (N) to noise temperature (T); and the familiar decibel-to-ratio conversions [6] for power gains and losses [7].

To calculate the system performance, start at the i. f. filter at the right-hand end of the block diagram, and work stage by stage towards the antenna. To avoid getting confused, work right through in terms of noise temperature alone, and then go back and convert all your system results to nf; similarly, calculate the cumulative front-end gain first in ratio terms, and then go back and convert all the results to decibels.

To set you going, the noise temperature at the input of the diplexer stage is given by [8]:

$$T = 225{\cdot}7K + 864{\cdot}5K/28{\cdot}2 = 256{\cdot}4K$$

This result is entered below the block diagram as shown in Fig 6. Now try calculating the noise temperatures at the inputs of all the rest of the stages

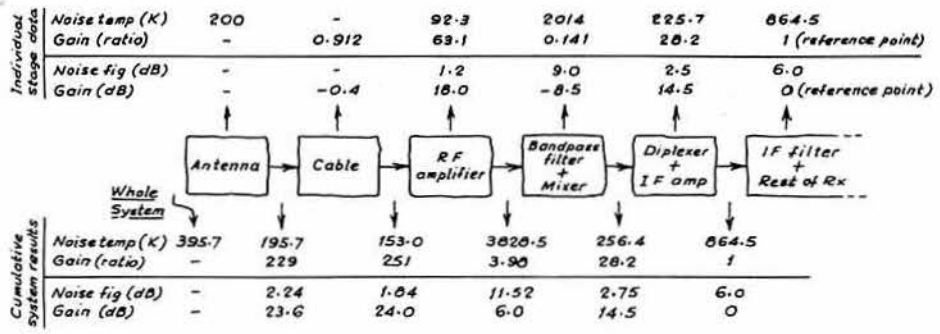
Individual stage data
Noise temp (K) 200 - 92·3 2014 225·7 864·5
Gain (ratio) - 0·912 63·1 0·141 28·2 1 (reference point)
Noise fig (dB) - - 1·2 9·0 2·5 6·0
Gain (dB) - −0·4 18·0 −8·5 14·5 0 (reference point)
Antenna
Cable
RF amplifier
Bandpass filter + Mixer
Diplexer + IF amp
IF filter + Rest of Rx
Whole System
Cumulative system results
Noise temp (K) 395·7 195·7 153·0 3828·5 256·4 864·5
Gain (ratio) - 229 251 3·90 28·2 1
Noise fig (dB) - 2·24 1·84 11·52 2·75 6·0
Gain (dB) - 23·6 24·0 6·0 14·5 0

Fig 6. Work sheet for noise temperature and gain analysis of front-end simplified from Fig 5

as far as the rf amplifier, and see that you get the same results as Fig 6. The noise temperature at the input of the rf amplifier is of course the noise temperature of the whole receiver.

To account for the loss in the cable, assume its physical temperature is 290K and use equation (4). Once again, write the result in the appropriate place below the block diagram, and check your answer against Fig 6. Finally, add on the antenna noise temperature of 200K to give the noise temperature of the entire system.

Having once worked through the system analysis, you can try changing the gains or noise temperatures of individual stages, or even the order of the stages, and see how the system responds. The computer program to be described later in the series makes it very easy to do a range of trial calculations, to obtain the desired noise temperature while using a minimum of gain. That is how I arrived at the system of Figs 5 and 6, in which the noise temperature at the top of the feeder cable almost exactly equals the assumed noise temperature of the antenna, meeting the design target recommended earlier.

For this month, I will leave you with two problems which do not need a computer program; in fact you will learn more about doing the calculations if you use a pocket calculator. You can solve the second problem by trial-and-error; the wrong answers are at least as informative as the right one!

1. If the rf amplifier in Fig 6 was at the masthead; ie, the rf amplifier and the cable are interchanged, what would be the new receiver noise temperature? *(Answer: 159·3K).*

2. This new noise temperature is now rather lower than we really need, and gives us the chance to get rid of some unnecessary rf gain. If we move the rf stage to the masthead, by how much can its gain be reduced? *(Answer: almost 2dB).*

Notes and references

[1] Thermal noise is not the only kind of noise in resistors; eg, additional noise arises if the resistor carries an externally-imposed current.

[2] *IEEE Standard Dictionary of Electrical and Electronic Engineering Terms*, 2nd edition. Wiley-Interscience.

[3] Imagine a system consisting of a lossy component between two amplifiers.

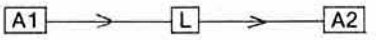

When calculating the nf of this combination, it is correct to add L (in decibels) to the nf of amplifier A2 (also in decibels), because this gives exactly the right answer for the nf of the combination, when checked against noise temperature using equations (4) and (5). But it is incorrect to regard L as simply reducing the gain of amplifier A1: this does not give quite the right answer for the nf of the combination, the error being due to having omitted the thermal noise generated in L itself.

[4] "136MHz/400MHz radio-sky maps", R E Taylor, *Proc IEEE* April 1973, p469.

[5] To relate microvolt sensitivity measurements to noise temperature:

(i) If the signal voltage is expressed as an emf, divide by 2 to obtain the equivalent pd in an impedance-matched system.

(ii) Calculate the signal power using $P = V^2/R$ where R is the nominal impedance of the measurement system (usually 50Ω).

(iii) Convert the signal : noise ratio from decibels into a power ratio. If the ratio is actually (signal + noise) : noise, convert it and subtract 1 to obtain the true s : n ratio.

(iv) Divide the signal power by the signal : noise ratio to obtain the *apparent* receiver noise power, referred to its input. Convert this to *apparent* noise temperature using equation (2) and assuming that the noise bandwidth B is the same as the −6dB i. f. bandwidth; this is usually a good approximation for ssb filters.

(v) To make the measurement, the receiver would have been connected to the signal generator via a resistive attenuator, which itself has a noise temperature of about 290K (ie room temperature). Subtract the attenuator's noise temperature from the apparent receiver noise temperature, to leave the *true* result. This correction is often forgotten, because it is unimportant when carrying out measurements on hf receivers which are not very sensitive; but it is vital for ultra-low-noise vhf and uhf receivers.

[6] If this is the first time you have ever found any use for decibels, the formulas for converting power gains or losses from ratios to decibels are:

$$G(dB) = 10 \log_{10} G \text{ (ratio)}$$

$$G \text{ (ratio)} = \text{antilog}_{10} (G(dB)/10)$$

Gains are ratios greater than 1, and are positive in decibels; losses are power ratios less than 1 (but always greater than 0), and are negative in decibels.

[7] Power transfers between stages are assumed to be in an impedance-matched system (conventionally 50Ω). If the system is not perfectly matched, the power gains should be replaced by "available gains" which allow for the mismatch losses.

[8] All the results in Fig 6 are quoted to several significant figures, so that you can be sure whether they agree with your own calculations. Normally you would round the answers off before writing them down.

TO BE CONTINUED IN THREE MORE PARTS

USING RESONANCE TO MEASURE CAPACITANCE

(Continued from page 263)

Ranges 6 to 1 should now be set-up (with C1 fully meshed and C2 half meshed) in that order, as shown in Fig 8, but substituting the new figures in column 2, ie 6 × swing of C1 on Range 7, 5 × swing of C1 on Range 6, 4 × swing of C1 on Range 5, and so on.

Extension of range

Although the upper limit of measurement is 980pF, it is perfectly feasible to measure much larger capacitors to a high degree of accuracy, by connecting the unknown in series with an accurately known standard, and measuring the capacitance of the series pair. The unknown may then be determined by use of the formula

$$Cx = \frac{C \times Ck}{Ck - C}$$

where Cx = capacitance of unknown capacitor
C = capacitance of series pair
Ck = capacitance of known standard

For example, an accurate 980pF capacitor might be selected (or made) by use of the instrument, and kept for use in this type of measurement. Suppose an unknown were connected in series with this, and the capacitance of the pair was measured and found to be, say, 892pF, then

$$Cx = \frac{892 \times 980}{980 - 892} = \frac{874,160}{88} = 9,933\text{pF}$$

There is a practical limitation to the use of this method for high-value unknowns, and this lies in the effect of inaccurate reading of the scale when measuring the capacitance of series pairs.

Uses

Many uses have been found for this instrument over the years. Experience has shown that the incidence of mismarked capacitors is higher than generally believed, and at my station all are tested before incorporation in a circuit.

Difficulty in achieving adequate temperature compensation in a vfo led me, some years ago, to investigate the characteristics of negative and zero temperature coefficient capacitors, by the use of a hair dryer. Several of the alleged ntc capacitors were found to have positive coefficients, and none of the zero coefficient capacitors had that characteristic. This was confirmed by changing the temperature in the opposite direction by means of a "freezer" aerosol.

An unfortunate attempt by an inexperienced listener to improve the ganging in his communications receiver by bending the adjusting vanes on the tuning capacitor, set a nasty problem. By the use of this instrument, and the expenditure of many hours, the four-gang capacitor was eventually returned to the maker's capacitance/rotation law.

The equipment has also been used successfully to locate the position of breaks in headphone and microphone leads, coaxial cables etc, by measuring the capacitance of the faulty core to screen or another conductor, from each end, and dividing the length *pro rata*.

Acknowledgement is made to the Jackson Brothers catalogue, for details of C28-142.

MODERN VHF/UHF FRONT-END DESIGN

Ian White, G3SEK*

PART 2. STRONG-SIGNAL PERFORMANCE

PART 1 of this article dealt with the performance of vhf/uhf receivers with the weak signals that we very much want to hear. This part deals with their performance with the strong signals that get in our way! Strong unwanted signals on the same frequency as a weak wanted signal are a fact of life, and the solutions are well known: good i.f. selectivity, a good sharp beam, and a good pair of ears with something in between them. This part of the article deals with a different type of problem, caused by strong signals which are *not* co-channel, yet still somehow manage to interfere with the wanted signal.

The general problem is often called front-end overload, though this blanket term covers several quite different effects. The three most important front-end overload effects are intermodulation, gain compression, and reciprocal mixing. There are many other overload effects, but these are of lesser practical importance in ssb/cw receivers. Designing and testing for good performance in these three main areas should take care of the other overload problems too.

Intermodulation

Intermodulation occurs when two or more signals mix together to produce extra signals that were not there before. Mixing between the local oscillator (lo) and the signal frequencies to produce the i.f. is the only kind of intermodulation that *ought* to occur in a receiver front-end. The other kinds of intermodulation that ought *not* to occur will produce spurious signals, and these may interfere with a weak wanted signal. It is important to realize that these spurious signals did not exist until they were generated *in the receiver*—but they cause just as much QRM as signals that arrive legitimately via the antenna!

The front-end has to cope with far more incoming signals than you ever hear. In a typical amateur radio vhf/uhf receiver the front-end is receiving all the signals present within a bandwidth of several megahertz, while you are hearing only those that get through the selective i.f. filter. At low signal levels all front-ends behave in a linear manner. This means that each incoming signal is processed separately, with no interactions between signals. But if incoming signals are very strong they can easily drive a poor front-end beyond its linear range. Then the signals begin to intermodulate; they mix together to produce new signals. So intermodulation is caused by non-linearity in the front-end.

If two signals on frequencies f1 and f2 mix, the first of the new frequencies to appear are (f1 + f2), (f1 − f2) and the second harmonics of each signal, 2 f1 (ie, f1 + f1) and 2 f2. These new frequencies are known as second-order intermodulation products (ips) because they were created by mixing *two* original "parent" frequencies. If the parent frequencies are close together, their second-order ips will be far removed in frequency, and can easily be filtered out [1].

If the two strong signals at frequencies f1 and f2 are increased in strength, another set of ips appears. These are the third-order products, so called because they involve mixing between three signals. These three signals could be totally independent, or the same effect can be generated by our two signals f1 and f2 on their own. Since these two frequencies can either add or subtract, the full range of possible third order ips includes:

the third harmonics (f1 + f1 + f1) and (f2 + f2 + f2);
the sum products (f1 + f1 + f2) and (f1 + f2 + f2); and
the difference products (f1 + f1 − f2) and (f2 + f2 − f1).

If both of the parent frequencies f1 and f2 are close to the wanted frequency, then the third-order sum products are all somewhere up in the third harmonic region and will not trouble us. But the two difference products containing a minus sign are somewhere close to their parent frequencies. For example, if f1 = 144·240MHz and f2 = 144·260MHz, their third-order ips fall on 144·220MHz (f1 + f1 − f2) and 144·280MHz (f2 + f2 − f1). Fig 7 shows how the two parent signals and their ips would appear on a spectrum analyser display. If the receiver front-end is allowed to generate these spurious signals, no practical kind of rf filtering can take them out, because they are too close to the frequencies you want to listen to. If one of these ips happens to fall on the frequency of a weak wanted signal, it will cause QRM. And if that *can* happen, then sure enough it *will*!

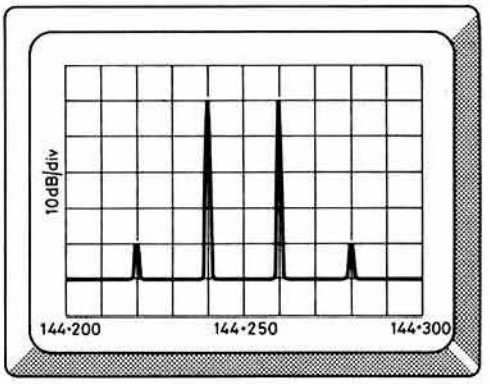

Fig 7. Spectrum analyser display of two equal carriers on 144·24 and 144·26MHz, with third-order ips at 144·22 and 144·28MHz

In the example above, the two strong signals were spaced 20kHz apart, and their third-order ips were evenly spaced 20kHz above and below them (Fig 7). As the strength of the two equal parent signals increases, the strengths of their third-order ips (also equal) come up three times as fast, increasing 3dB for every 1dB increase in the two parent signals. So the onset of intermodulation can seem quite sudden. Further increases in the strengths of the parent signals will produce new generations of higher-order ips. Thrashing through the algebra [2] shows that the odd-order (3rd, 5th, 7th etc) ips are the troublesome ones, because they can fall close to the parent signals. All these ips lie above and below their parents, at uniform frequency spacings (Fig 8). Fifth-order ips require stronger parent signals than third-order before they rise above the noise, but then they increase in strength five times as fast as the parent signals; and pro-rata for seventh- and higher-order products [2], which can thus appear very suddenly indeed if provoked by strong-enough parent signals.

Intermodulation requires at least two strong parent signals. If either one goes away, all their ip offspring will disappear. But relief will probably be short-lived, because the remaining strong signal can intermodulate with any other strong signal that comes along, to make a new set of third-order ips. And if the first signal comes back, making three in all, there will now be *four* (work it out) different sets of ips!

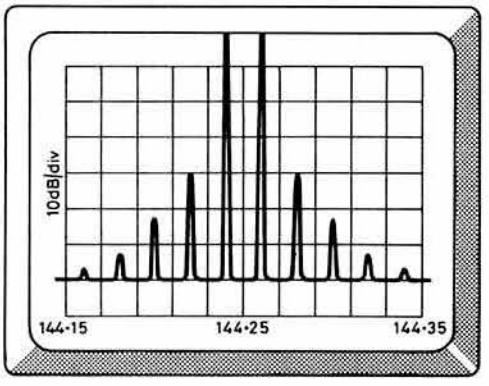

Fig 8. The same two carriers as in Fig 7, but increased in amplitude (off screen) to create severe intermodulation. Sweep expanded to show ips up to ninth order

*52 Abingdon Road, Drayton, Abingdon, Oxon OX14 4HP

So far we have been talking about single frequencies; ie, unmodulated carriers. There's worse to come. An ssb signal contains several different frequencies at once, and when these intermodulate in your receiver the effect is to slightly broaden the apparent bandwidth of the signal [3]. If two perfectly clean ssb signals intermodulate, then as well as each one seeming broader, their third-order ips together wipe out a further slice of spectrum, typically 20kHz. Remember, this is not the fault of the transmitters: the ips are generated in the receiver. With the low end of the band full of strong ssb and cw signals during a contest or an opening, intermodulation in the average commercial vhf/uhf receiver front-end can be ghastly. Listening on the band is like wading waist-deep through a shifting swamp of ips!

At this point I'd better inject a note of calm. *DON'T PANIC*, as the good book says [4]. It *is* possible to build vhf/uhf front-ends that have virtually no noticeable intermodulation, even at "open contest" signal levels. Compared with an ordinary receiver, the effect is startling: there are still plenty of strong signals around, but there are also spaces between them where the weak dx comes through in the clear.

Gain compression

Gain compression occurs when a strong incoming signal drives some stage in the front-end so hard that it can barely produce any more output. The gain of the "saturated" stage is thus reduced, taking with it all other signals, including the wanted one. In an extreme case, gain compression simply makes the receiver collapse in embarrassed silence whenever a strong signal comes on! Gain compression is easier to recognize if produced by a strong carrier than by ssb. When the carrier comes on, the background noise goes *down*, together with all other signals. If the carrier is keyed, the background noise appears as "negative" morse code, the dits and dahs being the silences in between. Note that gain compression requires only one strong signal, unlike intermodulation which requires at least two.

Reciprocal mixing

Reciprocal mixing is a large-signal effect caused by imperfections of the receiver's lo. Although reciprocal mixing has nothing to do with non-linearity of stages in the signal path, which causes gain compression and intermodulation, it is still a front-end problem. To understand how reciprocal mixing occurs, we need to take a short digression into the workings of oscillators and frequency synthesizers.

Starting with simple(?) oscillators, these in principle should produce only one frequency, corresponding to the resonance of the tuned circuit. Energy at this frequency goes round and round the oscillator's feedback loop, with the active device (usually a transistor) providing just enough gain to make up for circuit losses round the loop.

Unfortunately, oscillators never produce only one frequency: besides the main carrier and its harmonics, they produce energy at *all* other frequencies as well! Far away from the carrier frequency, there is a level background of noise due to the oscillator device and other resistive circuit losses. Other sources of wideband noise can include noise on the dc power supplies, and switching transients coupled in from nearby digital circuitry (especially of course in a synthesizer lo). All these sources of noise modulate the main carrier, and thus produce noise sidebands on either side of the carrier frequency. The closer to the carrier frequency, the higher the level of the noise sidebands, while at frequencies further away the noise sidebands sink below the level of broadband noise. Seen on a spectrum analyser, the result is a so-called "noise pedestal" standing out above the broadband noise and centred around the carrier (Fig 9).

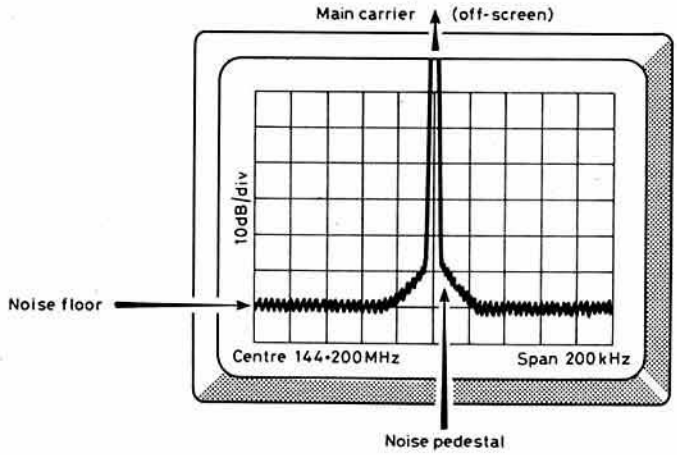

Fig 9. Narrowband noise spectrum of a carrier

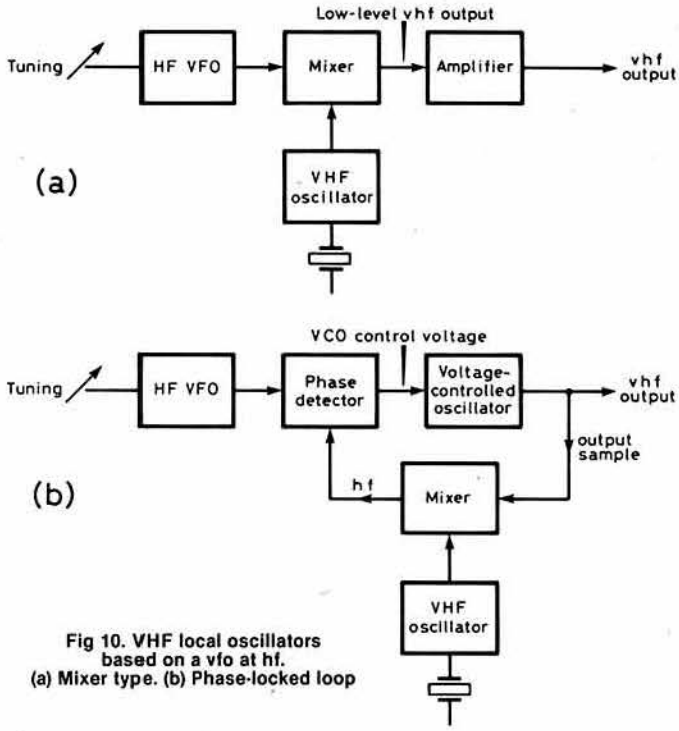

Fig 10. VHF local oscillators based on a vfo at hf. (a) Mixer type. (b) Phase-locked loop

Detailed analysis of the factors governing the levels of noise in oscillators is beyond the scope of this article. Suffice it to say that noise sidebands are better suppressed (relative to the carrier level) if the carrier level itself is high, and if the loaded Q of the frequency-determining tuned circuit is high. Oscillator design is far from simple, but it is quite easy to build a noisy oscillator by accident!

Because of the difficulties of obtaining adequate frequency stability in a vfo at vhf or uhf, transceiver los are usually either mixers or phase-locked loops (plls). Mixer vfos typically add the frequencies of an hf vfo and a vhf/uhf crystal oscillator (Fig 10(a)), and amplify the resulting output to a suitable level. Examples of transceivers using this technique are the old TS700 and the IC201. However, mixer vfos offer their own special opportunities for oscillator noise. The mixer cannot be run at too high a level without producing spurious products, and the carrier level at the output of the mixer can sink perilously close to the circuit noise. Any subsequent amplification of the low-level carrier also brings up the accompanying noise.

Other transceivers of similar vintage, such as the FT221 and FT225, use similar frequencies for the hf vfo and the vhf crystal oscillator, but add them by using a pll which is essentially a feedback control system. In Fig 10(b) the output signal is generated by a voltage-controlled oscillator whose frequency is also heterodyned down to hf using a vhf crystal-controlled oscillator. The phase of this hf signal is compared with that of the vfo, and any error voltage is fed back to control the vco. Tuning the hf vfo thus makes the vhf vco follow in step, at a constant frequency difference set by the crystal oscillator. That sounds fine: it produces the same output frequency as if the vfo and the crystal oscillator had been mixed directly, and the pll avoids the problems of spurious mixing products and noise in low-level circuitry. Or does it?

Phase-locked loop systems certainly *can* give very low noise sidebands if designed and constructed with care. But a bad pll can be very noisy indeed. In addition to all the noise problems that can beset ordinary oscillators, noise can easily be coupled onto the vco control line, leading to fm noise sidebands. Also, some designers have skimped on the normal precautions for generating a good-quality vco signal, relying instead on the phase-lock feedback loop to clean-up the vco signal for them. The result has been an lo with poor to bad levels of noise sidebands; a pity, because a little more care in design and layout could have produced a very "quiet" lo.

More recent vhf and uhf los are usually of the digitally-synthesized pll variety. These too *can* have good noise performance, but digital synthesis offers yet further opportunities to get things wrong! Parts of these synthesizers contain digital circuitry generating high-speed, high-level switching transients which can stray on to the vco control line unless careful precautions are taken. The frequencies fed to the phase detector in digital loops are quite low, and any leakthrough to the vco supply or control lines can cause modulation sidebands; in different systems these sidebands may

sound like carriers or like noise. It is very difficult to get rid of unwanted low-frequency modulation on the vco control signal, because filtering affects the ability of the loop to lock-up quickly and stay locked. The design of pll synthesizers thus requires some very careful trade-offs between various aspects of performance [5].

So much for our digression into oscillator noise. Now, what about reciprocal mixing?

Normally the mixer accepts a strong signal from the lo plus a weak wanted signal, and frequency-shifts the latter into the passband of the selective i.f. filter. But the mixer will do the same for any other pair of signals separated by the i.f. Reciprocal mixing gets its name because the roles of lo and signal are reversed: the strong signal is an unwanted one off-frequency, and the weak signal is part of the lo noise sidebands. For example, if you are listening to 144·250MHz with an i.f. filter centred on 10·700MHz, your lo is tuned to 133·550MHz (Fig 11(a)). The intention is that any signal appearing on 144·250MHz is mixed with the lo and delivered into the i.f. passband. Now if a strong signal appears on 144·26MHz, 10kHz away, it can mix with the lo noise sidebands around 133·56MHz, delivering that noise onto the 10·7MHz i.f. and raising the background noise level (Fig 11(b)). Note that the offset between the wanted frequency and the strong unwanted signal, 10kHz in our example, is the same as the offset between the lo carrier frequency and the interfering part of its noise sidebands. This means that the severity of reciprocal mixing will have the same frequency dependence as the lo noise spectrum. Reciprocal mixing between close-in strong signals and close-in lo noise is usually masked by leakage of the unwanted signal through the i.f. filter [6]. The main problems usually occur at the frequency offsets corresponding to the lo noise pedestal, though a particularly noisy lo may cause reciprocal mixing problems at frequencies all across its broadband noise floor.

On the air the effect of reciprocal mixing is opposite to that of gain compression. When a strong carrier appears, the background noise goes *up*. In a transceiver, the lo is used for both receive and transmit functions, so if the receiver suffers from reciprocal mixing, the transmitter will also radiate noise sidebands. It can be very hard to distinguish reciprocal mixing in your own receiver from the corresponding problem in someone else's transmitter. The only way to be sure is to test your own receiver independently, using a low-noise signal source.

Until fairly recently, reciprocal mixing was virtually unknown (or at least unrecognized) in vhf/uhf receivers for the amateur radio bands. This was partly because front-ends were even more vulnerable to intermodulation and gain compression than they are today, and also because los were generally quieter than their modern synthesized counterparts. Now that front-ends are getting better (slowly!), reciprocal mixing is becoming the limiting factor in strong-signal performance, and lo noise is coming under increasing scrutiny [5, 6, 7].

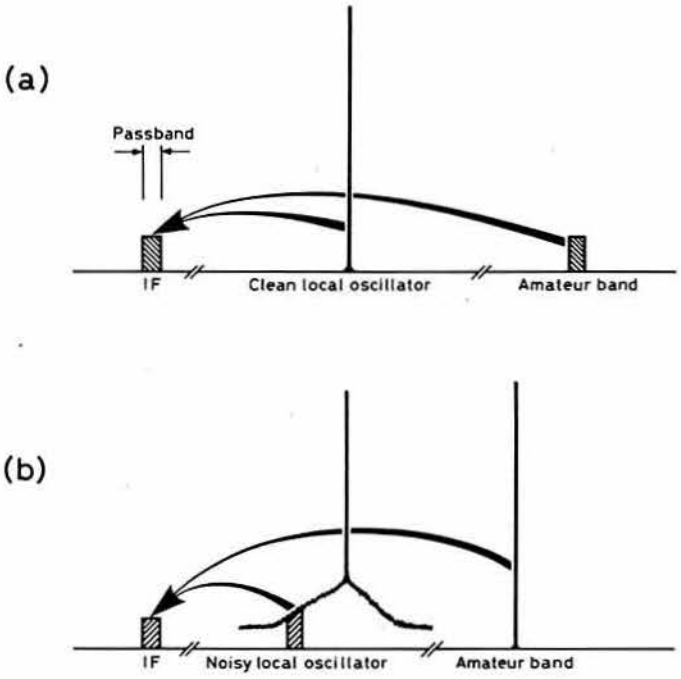

Fig 11. (a) Normal mixing. Local oscillator converts amateur-band signal into i.f. passband. (b) Reciprocal mixing. Strong amateur-band carrier converts local oscillator noise into i.f. passband

Dynamic ranges

We have seen that there are three main problems of front-end overload : intermodulation, gain compression and reciprocal mixing. They are all different, so each starts to be noticeable at a different level of strong unwanted signal. Two strong signals are needed to cause intermodulation, but only one to cause gain compression or reciprocal mixing. Intermodulation and gain compression are both due to some stage in the front-end being driven beyond its linear range, so keeping the gain in the front-end as low as possible helps cure both problems together.

However robust a front-end may be, it can still be overloaded by a strong-enough signal. Our objective is not to "cure" overload, but to make sure that overload hardly ever happens, even at the strongest amateur-band signal levels. For all practical purposes intermodulation is non-existent if the ips never rise significantly above the background level of system noise. Similarly, the receiver can be considered immune to gain compression or reciprocal mixing if these effects never significantly change the system noise level.

This principle, that what you don't hear doesn't trouble you, leads directly to the idea of spurious-free dynamic range (sfdr). The bottom end of the dynamic range is the power level of the weakest audible signal. The top end is the power level of the strongest off-frequency signal that the receiver can tolerate without the particular overload effect becoming noticeable. And the sfdr is the difference in decibels between these two power levels. Since each overload effect appears at a different level of strong signal, the receiver possesses a different sfdr for each.

The idea of spurious-free dynamic range relates directly to the things you hear when using a receiver. However, it needs a certain amount of further definition to turn it into something measurable. The power level of the "weakest audible signal" is arbitrarily defined as equalling the receiver's effective thermal noise power, $kT_{RX}B$ (recall Part 1). Other names for this power level are minimum discernible signal (mds) or noise floor. Formal definitions for the points at which the various overload effects become significant have been proposed by Wes Hayward, W7ZOI, and these lead directly to methods of measuring the various types of sfdr [6, 8]. The W7ZOI dynamic range tests are based firmly on reality—what you measure is also what you hear on the air—so it is not surprising that they have become the standard for hf equipment reviews in *QST* and *Rad Com*. However, the test methods described by Hayward were developed for hf receivers, and need some further development for use at vhf and uhf; but that's another story.

Notes and references

[1] If signals on a very wide range of frequencies are allowed to enter the front-end, and they then intermodulate, their second-order ips can fall close to the wanted frequency. This can occur if there are two extremely strong fm or tv broadcast transmissions on frequencies about 144MHz apart. Second-order ips from these two wideband signals can cover the 144–146MHz amateur band. Fortunately the cure is simple: the unwanted parent signals are so far in frequency from 144MHz that they can easily be kept out of the receiver by filters.

[2] Chapter 6 of the excellent *Introduction to Radio Frequency Design*, W A Hayward (Prentice-Hall Inc, Englewood Cliffs, New Jersey, 1982, ISBN 0-13-494021-0) gives a clear account of the algebra of intermodulation.

[3] Intermodulation occurs in transmitters as well as receivers, and the two can be very difficult to distinguish. Intermodulation in ssb transmitters is responsible for the reintroduction of signals on the opposite sideband (which had earlier been filtered out), and hence for broadening of the signal. The same broadening is heard if the intermodulation occurs in the receiver. Being of fairly low order, these ips are quite close to the main signal and their level usually drops quite rapidly as you tune away; they sound like unintelligible "mush". Spiky-sounding, impulsive "splatter" or "buckshot", often spreading over several tens of kilohertz, is due to very-high-order intermodulation which is much more likely to be generated in overdriven transmitters than in overloaded receivers—but that's also another story!

[4] *The Hitch-Hiker's Guide to the Galaxy* (Megadodo Publications, Ursa Minor and Surbiton, various dates).

[5] *Introduction to Radio Frequency Design* (see note 2), chapter 7. *Frequency Synthesis by Phase Lock*, W F Egan (John Wiley & Sons, New York, 1981).

[6] *Solid State Design for the Radio Amateur*, Wes Hayward, W7ZOI, and Doug DeMaw, W1FB, chapter 6. (ARRL).

[7] See [5] and [6] and also;
"Dynamic Range, Intermodulation and Phase Noise", Peter Chadwick, G3RZP. *Rad Com* March 1984, p221.
"Modern Receivers and Transceivers: What Ails Them?", Wes Hayward, W7ZOI, and Doug DeMaw, W1FB. *QST* January 1983, p11.

These references, together with equipment reviews in various magazines and frequent items in *Technical Topics*, show the increasing attention being paid to lo noise in receivers and transceivers. However, most commercial transceivers fail to meet our expectations. As Pat Hawker and others have pointed out, the introduction of synthesizers was mostly to avoid the labour costs of building and calibrating traditional LC vfos. For dx working on the hf, vhf or uhf amateur bands, the scanning and memory facilities of synthesizers are a poor recompense for the attendant lo noise, which in many transceivers impairs the most basic functions of receiving and transmitting.

[8] "Defining and Measuring receiver Dynamic Range", Wes Hayward, W7ZOI. *QST* July 1975, p15. Although the same information is incorporated in Hayward's *Introduction to Radio Frequency Design* and in *Solid State Design for the Radio Amateur*, the original article is more detailed and well worth reading.

TO BE CONTINUED

4-2-70

*by Ken Willis, G8VR**

THE 1985 NATIONAL VHF CONVENTION, held at Sandown Park Racecourse on 23 March, proved to be bigger and better than ever. More people passed through the turnstiles than at any previous such event, while the number of traders exhibiting was considerably up compared with earlier years. The interest in a convention such as this stems mainly from the fact that current licence conditions make it necessary for Class B operators to confine their activities to the vhf/uhf bands, whereas years ago vhf was a specialist activity, forcing those few wishing to use these frequencies to build their own equipment since nothing much was available from commercial sources. The Japanese "black-box" changed all that, virtually overnight, so today's vhf operators are not all cast in the same mould, but instead display an interest in many facets of vhf/uhf communication, not simply the quest for better front-ends, bigger antennas and the like which still remain the domain of the so-called weak-signal operators. To many vhf operators today, the fact that their activities take place at vhf is perhaps co-incidental, since they could do much of what they wanted on any other frequency.

The convention will, according to custom, be fully reported in a later issue of *Radio Communication* with joint-authors covering the three lecture streams which formed the main afternoon programme. It was good to meet in person so many enthusiasts, several having travelled really long distances to attend the event. When Sandown Park was chosen as an alternative to the old "Winning Post" at Twickenham, many thought that the new venue was altogether too large, but the crush on the floor this year, which at times made it difficult to move around at all, suggests that we shall even grow out of this location eventually.

VHF/UHF Newsletter

During one of the lecture sessions at Sandown Park, the VHF Committee sat in forum to respond to questions from the floor. Some discussion arose on the relative roles of the recently-initiated *VHF/UHF Newsletter* and *4-2-70*. One speaker went so far as to suggest that the emergence of the newsletter had resulted in copy being "filched" from *4-2-70*. This is by no means the case. The newsletter is aimed at providing a quick response to vhf matters too urgent to await the normal periods required to get a journal such as *Radio Communication* from rough copy to final print stage. In addition, *4-2-70* (and its predecessors under various titles) has for many years not been regarded as a vehicle for technical information of the sort requiring circuit diagrams, component details and the like, since such matters are typically covered by *Technical Topics* or by specialist articles.

4-2-70 aims to provide news and views of a general nature on all topics related to frequencies between 50 and 500MHz. It is *your* feature and must always reflect the input received from readers over a wide range of subjects where vhf or uhf is the common denominator. It tries to strike a balance between the various interests of readers, be they repeaters, weak-signal operation, sstv, rtty, satellites etc, and if at any time the emphasis on any one topic appears too great, it simply reflects a large input from readers related to those matters at the time. For example, very little has been happening in recent months to make the dx-minded operators get very excited. On the other hand there have been some interesting developments in the repeater field, with some engineering information being received which indicates the level of sophistication which some of these "machines" now possess. Packet radio is a subject being very much talked about, while the use of the micro in our hobby becomes ever more important. 50MHz has sometimes dominated the available space because it is a very active band now (during certain hours) and likely to be one which we will all be able to use at sometime in the future. So all these things I regard as meat for *4-2-70*, though I may not myself have a personal interest in all of the topics reported to me.

Returning now to the newsletter, this will, from time to time, include circuits and technical information of a detailed nature. It will also include material of an operational nature, mainly of interest to weak-signal operators (including eme and ms), and stop-press type information which can be distributed more quickly through the newsletter because of its simpler format and method of preparation. The editor of the newsletter does not have access to information sent to *4-2-70*, nor do I see his copy prior to publication. Indeed, Dave Butler and I go to extreme lengths to avoid seeing one another's copy, since this might inhibit either or both of us from including a specific item in order to avoid duplication. Duplication will sometimes occur, since there are sources of information used by technical writers in a given field which are distributed "at large". You will often notice duplication between the various magazines and journals in our subject field for this reason, but to omit something of interest simply because it might appear elsewhere would not be a sensible way of keeping readers informed.

4-2-70 is *your* feature, as is the *VHF/UHF Newsletter* for those who subscribe to it. It is in your own interest to tell us what you want or expect from these publications, and to back up your demands by contributing, whenever possible, items of interest which you think deserve a more general dissemination.

Activity in some rare Spanish squares

The *SM6EOC/SM6AFH 2 Metre News-sheet* gives details of stations active in some rare Spanish squares. Olle and Thomas seem to prefer the old notation, so it is presented here just as they stated it. From XY square, EA7APD and EA7BVD are both active on tropo and also monitoring for sporadic-E openings. From ZA, EB4IM is similarly active, while in YY, EA7CGH and ED7GEL will be stations who will be on the look-out for dx contacts. For those who prefer meteor scatter, EA4CVS in XZ is equipped for ssb-ms as well as tropo and Es monitoring. Skeds can be arranged by telephoning Adolfo on 927-532331, or write to him QTHR (in International callbook). Having recently paid a visit to some amateurs in EA5, I can testify to the tremendous problems which face the vhf operator in those parts due to high mountains which are a feature of the area. "Fives" are rare on vhf, but from ZX square, EA5BYS and EA5BSH are both active on tropo and looking this year for Es contacts. Meteor scatter can be useful in getting over some of those high peaks, and therefore the presence of EA5EMM and EB5EHX in ZZ square, both active on cw ms (maximum 1,000lpm) is good news for the squares hunters. Skeds may be made on the 14MHz vhf net.

Olle also says that, at the present time, beacon EA1VHF is off the air while some problems of harmonic radiation from it are cleaned up.

Repeater news

It is understood that the problems associated with the widespread coverage of the new IoM repeater GB3GD (R1) have been largely overcome. From its vantage point on top of Snaefell, some 2,036ft asl, the signal put out was enormous and heard over considerable distances, causing some co-channel interference.

*11 Old Downs, Hartley, Kent DA3 7AA

MODERN VHF/UHF FRONT-END DESIGN

Ian White, G3SEK*

PART 3. INTERMODULATION ANALYSIS

THIRD-ORDER INTERMODULATION is one of the most common and annoying front-end overload effects. However, it can be analysed, and we can use the analysis to design front-ends. In general, a front-end designed to be resistant to third-order intermodulation will also be resistant to intermodulation at higher orders, and probably also to gain compression. So let's take a closer look at third-order intermodulation.

In Part 2, I explained that third-order intermodulation is an effect caused by two strong parent signals, which drive the front-end beyond its linear range and produce spurious signals called intermodulation products (ips). I also mentioned that third-order ips increase in amplitude three times as quickly as the pair of equal parent signals. Fig 12 plots the power levels of the signals coming out of a "black box" (eg, an amplifier or a mixer) against the input power of either one of the two parent signals. With no input signal, the output is merely noise at the noise-floor power level. At low input levels (point A), the only output signal is the amplified parent, increasing at 1dB/decibel of increase at the input. At higher input levels, the third-order ips appear out of the noise floor (point B), and then increase at 3dB/dB. Ultimately, it looks as if the level of the ips would catch up with the main output signal. The point where they would do so is the third-order intercept point (I).

The strength of the intercept-point concept lies in its simplicity. Its weakness is that intermodulation intercept points do not really exist. Before the ips can catch up with the parent signals, the amplifier or mixer goes into gain compression and its output level flattens off (Fig 12). Thus the intercept point cannot be measured directly; it can only be determined by a theoretical extrapolation of measurements made at lower signal levels. There is no guarantee that real systems will behave in the obliging manner

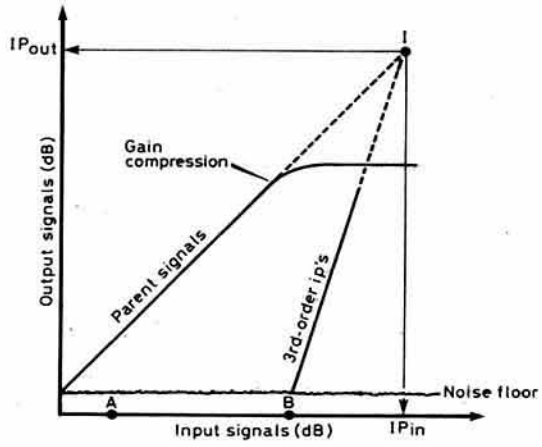

Fig 12. Third-order intermodulation plot, showing intercept point I and corresponding input and output levels IP in and IP out

*52 Abingdon Road, Drayton, Abingdon, Oxon OX14 4HP

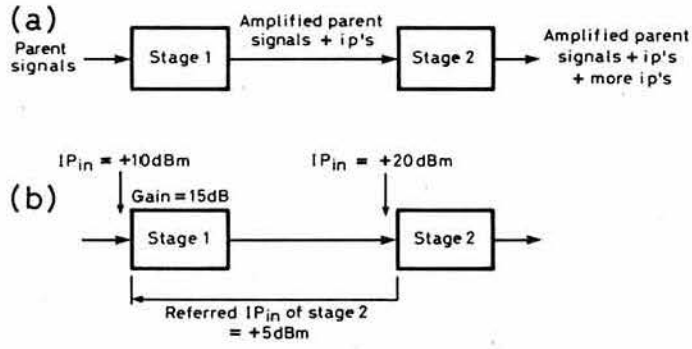

Fig 13. (a) Intermodulation in cascaded stages. (b) Input intercept of Stage 2 is referred to input of Stage 1 by subtracting Stage 1 gain

that the simple theory demands; sometimes they do not. However, agreement between theoretical predictions and *accurate* measurements is generally pretty good, and one thing is for sure: a system designed with the aid of the intercept-point theory will be a lot closer to optimum than a system that has not been designed at all!

In Fig 12 the intercept point (I) can be specified in two ways, by reference to the corresponding input power, IP_{IN}, or to the output power, IP_{OUT}. Either is valid (accepting the fact that both are purely theoretical), and the two differ simply by the small-signal power gain of the stage in question. If an amplifier with a power gain of 19dB has an output intercept of 30dBm (decibels over 1mW), its input intercept is (30 – 19) dBm; ie, 11dBm. Note that the output intercept of an amplifier *looks* better than the input intercept, so beware of anyone who avoids mentioning which intercept he's quoting—especially if he's trying to sell you something! On the other hand, if a diode ring mixer with a *loss* of 6dB has an output intercept of 10dBm, its input intercept is 6dB *greater;* ie, 16dBm. The third-order intermodulation intercepts of individual amplifiers, mixers etc can usually be estimated either from manufacturers' design data or by some simple rules-of-thumb, as we shall see later. It then becomes possible to predict the intermodulation performance of a multi-stage front-end, in the same way as we predicted its small-signal performance in Part 1.

So far we have been thinking about intermodulation in single stages. Now, what about intermodulation in cascaded stages that make up a complete front-end? If two equally-strong signals are fed into the input of a two-stage system (Fig 13(a)), the first stage will produce an amplified output, plus some internally-generated ips. The next stage will accept all these as inputs, and as well as amplifying both signals and ips it will produce yet more ips of its own making. The two sets of ips will be on the same frequencies (having the same parent frequencies), and will usually add in phase at the output of the second stage [1]. If the input intermodulation intercepts of stages 1 and 2 in Fig 13(a) are $IP_{IN}1$ and $IP_{IN}2$, the system input intercept is given by:

$$1/IP_{IN}sys = 1/IP_{IN}1 + 1/(IP_{IN}2 / G1) \quad \ldots\ldots(6)$$

Let's look at this formula more closely. It may seem vaguely familar, because it is rather like two formulas you have seen before. One is the formula for adding noise temperatures in cascaded stages:

$$T_{SYS} = T1 + (T2/G1).$$

That was equation (2) in Part 1, where I mentioned that the action of dividing T2 by G1 was called "referring T2 to the input", so that it could be added directly to the noise temperature T1 of the first stage. In just the same way we are referring $IP_{IN}2$ to the input when we divide it by G1 in equation (6). The other familiar aspect of (6) is that when the two intermodulation intercepts have both been referred to the input, they add together just like resistors in parallel. As you know, the resistance of a parallel combination is always less than the lowest of the individual resistances. It is exactly the same with intermodulation intercepts: the input intercept of the system is lower than the lowest of the individual stage intercepts—all intercepts having first been referred to the system input. Unlike the chain in the proverb, this one is *weaker* than its weakest link.

Time to try some numbers. Fig 13(b) puts typical values to the general system of Fig 13(a). Using equation (6)

$$1/IP_{IN}sys = 1/(10dBm) + 1/(20dBm - 15dB)$$

This equation cannot be evaluated as it stands. The intercept of the second stage can be referred to the system input simply by subtracting the 15dB gain of the first stage, but power levels in decibels relative to 1mW must be converted to milliwatts before they can be combined. So, more correctly,

$$1/IP_{IN}sys = 1/(10mW) + 1/(3 \cdot 16mW),$$

so $IP_{IN}sys = 2 \cdot 40mW$ or $3 \cdot 80dBm$. Note that the result is lower than either of the individual stage intercepts.

What happens if one stage has a much higher intercept than the other

(both of course having been referred to the system input)? If the second stage is the more "robust", the system input intercept still cannot be greater than the 10dBm limit imposed by the first stage. Conversely, no matter how robust the first stage, IP_{IN}sys cannot be better than 5dBm owing to intermodulation in the second stage. Note also how the effective performance of the second stage is made worse by the gain ahead of it in the first stage. Time after time, we keep coming back to the need to keep the system gain down.

Worksheet methods

In part 1 we used a worksheet (Fig 6) to analyse the noise temperatures and noise figures of a 144MHz front-end. Now we will use the same worksheet technique to analyse its third-order intermodulation performance. Unlike the noise temperature analysis, which would have been incomplete without considering noise from the antenna, the intermodulation analysis can be confined to the receiver alone. Fig 14 shows the front-end with the same stage gains as before, and with typical intermodulation intercepts for each stage. Later I will suggest some ways of estimating these intercepts.

Each intercept is referred to the input of its stage, so the first thing to do is to refer all intercepts to the receiver input. This is done by subtracting the cumulative gain up to that point, remembering that we are working in decibels: check Fig 14. To combine the stage intercepts into a system value, convert everything into milliwatts, and add all the values as if they were resistors in parallel—go on, get out your calculator and do it! The answer you should get is 0·299mW, which converts to −5·23dBm as shown in Fig 14.

So the third-order input intermodulation intercept of this front-end is −5·23dBm. What else does the analysis tell us? As soon as we referred each of the stage intercepts to the input, we could see straight away which are the vulnerable stages. The mixer is the most vulnerable, for its referred intercept (−3dBm) is the lowest of all. Next comes the rf stage (0dBm) and after that the i.f. filter. Thus the first step of improvement would be to change the diode mixer to one with a higher input intercept; but that would only improve the system intercept a little, because of the 0dBm intercept of the rf stage.

Another thing the analysis tells us is the predicted dynamic range for third-order intermodulation. Recalling from Part 1 that the receiver noise temperature was 153K, its noise floor in a typical ssb bandwidth of 2·5kHz is −142·8dBm (from equation (1)—check for yourself). The so-called spurious-free dynamic range (sfdr), also known in this context as the intermodulation-free dynamic range or the two-tone dynamic range, is given by:

sfdr = (input intercept—noise floor) × 2/3 ..... (7)

everything being measured in decibels; ie,

(−5·23dBm − −142·8dBm) × 2/3 = 91·7dB.

Measurements on front-ends like that in Fig 14 agree quite well with our predictions, and so does experience on the bands. Assuming a decently-calibrated S-meter (a big assumption!), with 6dB per S-point and S0 at the noise floor, absolutely no third-order intermodulation should be noticeable until signals reach almost full scale; ie, about S9+40dB. That is the kind of performance you should expect from a modern front-end.

Estimating intermodulation intercepts

If you went ahead and built the whole receiver, had enough test equipment, and also had the knowledge to use it, you could make measurements without having to do all this paper analysis. But that doesn't measure up to my idea of the word "design". The reasons for doing the paper design and analysis are to avoid wasting your time building a front-end that you could have predicted to be inferior, and also to substitute for test equipment which we amateurs don't possess. Of course, a fully-professional approach to front-end development requires both analysis and measurement.

The whole concept of third-order intermodulation analysis rests on three assumptions: that intermodulation in each stage can be characterized by an intercept; that intercepts add up, as shown earlier; and that you know what values to assume. I have already dealt with the first two assumptions. They may not be totally accurate but they're far better than nothing at all. So now we need some information on values; and if manufacturers' datasheets shed no light we can estimate the intermodulation intercepts for ourselves.

Intermodulation intercepts in amplifiers can be estimated by thinking first about gain compression. A Class A amplifier without negative feedback can be considered simply as a controllable resistor in series with a load impedance. In bipolar transistor amplifiers the control is via base current; control in fets is via gate-source voltage, in valves via grid-cathode voltage. The same principles apply to all devices, so for simplicity let's talk in terms of bipolar transistors. The output of the amplifier will be limited either if the device is instanteously turned hard on; ie, the collector-emitter voltage V_{CE} swings right down to zero (voltage limiting); or alternatively if the device is instaneously turned completely off and no current flows (current limiting). Any further increase in drive will result in severe compression. Hayward shows [2] that the maximum output powers in these two cases are approximately:

$V_{CE}^2/2R_L$ for voltage limiting, and $I_B^2 R_L/2$ for current limiting,

where R_L is the load resistance, and I_B is the Class A bias current.

The lower of these two power levels is the maximum undistorted power output of the amplifier. For a bipolar amplifier biased to 20mA at 6V, and terminated in 50Ω, current limiting is the operative factor and occurs at an output power of 10mW. The same upper limit would apply to *any* device, including a GaAsfet or even a valve, biased to the same I_B and terminated in 50Ω. For an old-fashioned bipolar device giving its lowest noise figure at an I_B of only 2mA, current limiting would occur at a mere 0·1mW output —no wonder they had a bad reputation for overload! In these applications, voltage limiting hardly comes into the picture at all, though it may do so for fets or valves terminated in high rf load resistances [3].

Subtracting typical stage gains of 15-20dB or more for modern low-noise devices, the input powers at the compression point will be −10dBm or worse. As a rough rule of thumb, the third-order input intercept of a *good* bipolar amplifier is about 15dB above its compression point [2], and the same may be true of some fets. Badly-biased or non-linear devices may be much worse! Thus the input intercepts of present-day low-noise amplifiers are most unlikely to be above +5dBm, and 0dBm is far more typical [4]. Negative feedback gives some improvement in linearity, and also gets rid of surplus gain, so it can be a highly effective means of breaking out of this "intermodulation trap". However, negative feedback does tend to detract from low-noise performance, and even with heavy feedback the fundamental limitations on *output* voltage or current swings still apply.

In Fig 14 the input intercept of the rf amplifier is assumed to be 0dBm in accordance with the above estimates, while the value of +20dBm assumed for the i.f. amplifier is typical of a feedback amplifier using a bipolar transistor designed for linear operation and biased to about 50mA.

Measurement on diode-ring mixers show apparent variations in third-order input intercepts so the simple theory does not work too well for these mixers. Some of the reasons are not too hard to see. All the power coming out of one of these passive mixers must be supplied either from the incoming signal or from the lo. Thus there is competition for the available power, and the level of each wanted or spurious signal depends on all the other levels as well as on the non-linearities of the mixer. However, as a rough guide the third-order input intercept of a *properly terminated* passive ring mixer is about 5dB above the lo power level [5]. DeMaw and Collins have examined the intermodulation performance of some other types of mixers at hf [6]. Mixers with gain are most likely to limit at their outputs, and can be assessed like amplifiers.

It is not safe to assume that *any* passive circuit device is immune to intermodulation without careful examination of how it works. Resistors are

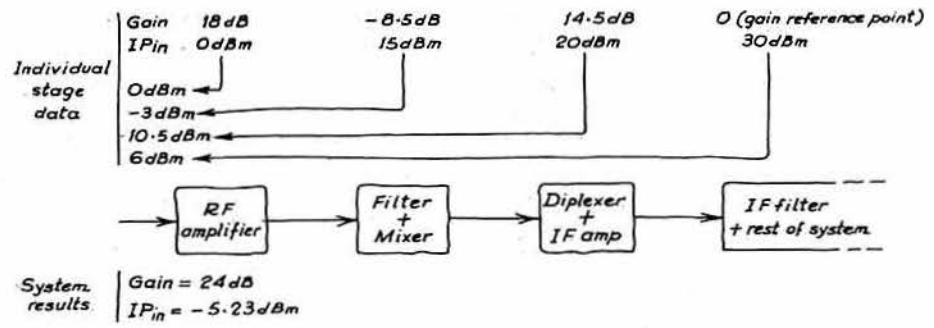

Fig 14. Intermodulation analysis of receiver from Part 1 of this article

a fairly safe bet, according to Ohm's Opinion, and so are most capacitors. But variable-capacitance diodes will intermodulate impressively if the instantaneous signal voltage nearly cancels out the reverse bias. Similar problems can occur with all kinds of diodes (including the "parasitic" diodes on transistor chips) so diodes should be firmly biased either on or off, and varicaps should be avoided. Any kind of ferrite or dust-iron cored device can cause intermodulation unless the magnetic material is operated a long way below saturation. Peter Chadwick, G3RZP, has pointed out the possibilities of intermodulation even in such apparently harmless devices as crystal filters [7], and we saw in Fig 14 that this can indeed be significant in high-performance front-ends.

Finally, remember that all these methods of estimating intermodulation performance are approximate, and that they generally apply to idealized devices of their type. The performance of badly-chosen or badly-used devices will be *worse!*

The TCALC program

TCALC is a computer program to analyse the noise and intermodulation performance of receiver front-ends. It works on exactly the same principles as the noise analysis in Part 1 and the intermodulation analysis described earlier. Stages are numbered backwards towards the antenna, starting with Stage 1 which is always the "rest of the system". The program works interactively, stage by stage. It first asks you what kind of stage is coming next, (eg, an amplifier or an attenuator) and then prompts you to enter appropriate data. TCALC then calculates the noise and intermodulation performance of the system, referred to the input of that stage, and then asks you about the next stage.

A particular feature of TCALC is the ease of repeating a calculation. You can change the input data if you want, or a single keystroke recalls the old data for any stage you want to leave unaltered. Thus you can quickly optimize the system performance by comparing new and old results.

Two versions are available, one written for the BBC micro and the other in Microsoft Basic-80, the nearest thing to a common Basic dialect. Both versions are written for ease of understanding rather than economy or speed, for they occupy little memory and run very quickly anyway. Of the two, the Microsoft version is recommended for easier translation into other dialects of Basic.

A listing of the Microsoft Basic version of TCALC appears in the June 1985 issue of the RSGB *VHF/UHF Newsletter*. You can also obtain either the BBC Basic or the Microsoft Basic listing by sending an A4 sae to *VHF/UHF Newsletter,* PO Box 73, Hereford HR2 9EW. Please don't forget to state which version you want.

Noise analysis

The program starts by telling you that each stage can be called either an AMPlifier, a MIXer, an ATTenuator, a CABle or an ANTenna. Each label is recognised by its first three capital letters, and treated accordingly.

Inside the program, an AMPlifier stage is treated the same as a MIXer, the different labels being purely for your own convenience. The program asks you for the stage's noise temperature or noise figure (nf); for convenience a value less than 20 is assumed to be a nf (in decibels), while any greater value is assumed to be a noise temperature (in Kelvins). AMP/MIX stages also have gain, so TCALC asks for an input in decibels; negative gain values (losses) are of course allowable. Internally the program works in terms of noise temperatures and ratio gains/losses, as I recommended in Part 1. TCALC automatically treats the first stage, the "rest of the system", as an AMP/MIX type and sets its gain to 0dB as we did in the worksheet example.

ATTenuators and CABles are the same as far as the TCALC is concerned, and the only input requirement is the loss in decibels. CAB or ATT should also be used for other devices that introduce attenuation, such as filters. The program assumes that the attenuator is at a physical temperature of 290K in order to calculate the noise contribution from its resistive losses.

The only input requirement for the ANTenna is its noise temperature. TCALC recognizes that the antenna is always the final stage, though if you want to finish the calculation earlier you can input END. The program finally displays the system results and asks if you want to repeat the calculation.

On repeating the noise analysis calculation, you are once again prompted stage by stage for input. If you do not want to change the characteristics of any stage, you can re-use the old data by simply pressing the "return" key. To change the input data for any stage you re-specify its type as AMP, MIX etc, and TCALC then asks for your new input data.

Fig 15 is the printed record of a noise analysis using TCALC, corresponding to the worksheet in Part 1 (Fig 6). Note that for the final stage, the antenna, only the total noise temperature is displayed (because noise figure and cumulative gain have no meaning in this context).

Intermodulation analysis

The input intermodulation intercept of each stage needs to be referred to the system input, before all can be combined using equation (6) to give the input intercept for the whole system. For TCALC to do this, it must already know the gain of every stage in the front-end, so you have to start with a noise analysis which gives the program the gain information it needs. In the

```
Stages are numbered backwards towards antenna
'Stage 1' always represents the rest of the system: specify its noise figure
or noise temperature.
For all subsequent stages you are asked to specify the type.
  The following types of stages are recognised by their first 3 CAPITAL letters:
AMPlifier or MIXer, CABle or ATTenuator, ANTenna, END.

STAGE PARAMETERS                          SYSTEM PARAMETERS
                                    T(K)        NF(dB)      G(dB)
Stage 1 (rest of system)
T(>20K) or NF(<20dB): 6.0
                                    864.5       6           0
Stage 2 Type: AMP
T(>20K) or NF(<20dB): 2.5
Gain (dB): 14.5
                                    256.4       2.75        14.5
Stage 3 Type: MIX
T(>20K) or NF(<20dB): 9
Gain (dB): -8.5
                                    3828.5      11.52       6
Stage 4 Type: AMP
T(>20K) or NF(<20dB): 1.2
Gain (dB): 18.0
                                    153         1.84        24
Stage 5 Type: CAB
Loss (dB): 0.4
                                    195.7       2.24        23.6
Stage 6 Type: ANT
Ant noise temp (K): 200
                                    395.7
================================================================================
Repeat, Intermods or Quit (R/I/Q)?
```

Fig 15. Printer output of noise analysis using TCALC, corresponding to the Worksheet (Fig 6 in Part 1)

```
THIRD-ORDER INTERMODULATION ANALYSIS

INPUTS
3rd order input intercepts for 'rest of system' and AMP/MIX stages.
Enter values in dBm, or RETURN to input a high value.

OUTPUTS
For each stage: 3rd order input intercept, referred to system input.
 For whole system: input intercept (dBm), noise floor (dBm), dynamic range (dB).

STAGE PARAMETERS                          SYSTEM PARAMETERS
                                  T(K)    NF(dB)    G(dB)    IPi3(dBm)
Rest of system ip3 (dBm): 30
                                  864.5   6         0        6
Stage 2 (AMP) ip3 (dBm): 20
                                  256.4   2.75      14.5     10.5
Stage 3 (MIX) ip3 (dBm): 15
                                  3828.5  11.52     6        -3
Stage 4 (AMP) ip3 (dBm): 0.0
                                  153     1.84      24       0

Whole receiver:                                               -5.2
Noise floor (2.5kHz) = -142.8dBm
Dynamic range = 91.7dB
================================================================================
Repeat, Intermods or Quit (R/I/Q)?
```

Fig 16. Printer output of TCALC intermodulation analysis corresponding to Fig 14

intermodulation analysis, TCALC steps through stage by stage, displaying the existing data on stage type and cumulative noise temperature, nf and gain. Only if the stage is an AMP/MIX will you be asked for its input intercept. CAB/ATT stages are assumed to be intermodulation-free [8], and the same can be assumed for AMP/MIX stages by simply pressing "return" instead of entering an intercept value in dBm. The program then displays the input intercept of each stage, referred to the system input. TCALC recognizes when it has arrived at the system input, and then displays the system input intercept, the noise floor in decibels relative to 1mW and the intermodulation-free dynamic range in decibels.

Fig 16 is the printed record of an intermodulation analysis corresponding to the worksheet (Fig 14). There are two ways of entering the +30dBm intermodulation intercept of the i.f. filter: the method used in Fig 14 is to absorb the filter into the "rest of the system" and assign a +30dBm intercept to the latter; another more general way is described in [8]. Note also that for the intermodulation analysis the chain of stages in Fig 16 has been cut short at the receiver input.

Using TCALC

After typing TCALC into your own micro, translating the Basic where necessary, use Figs 15 and 16 to check that it works properly. Also make sure that the results of previous calculations are accurately recalled. The general objective of the screen layout is to keep inputs on the left-hand side and results on the right, taking advantage of an 80-column display. If your display is narrower you may have to make some compromises, and old results will disappear more rapidly off the top of the screen.

TCALC is not intended to make noise and intermodulation analyses more accurate: it is intended to make them more convenient. With old-fashioned front-ends, with their vast excesses of unnecessary gain, you could get away with comfortable generalizations about the effects of individual stages on system performance. But those generalizations are not true in modern front-end design. Today we are trying to keep front-end gain as low as possible, and minor changes to a single stage can affect the performance of the whole system. Even so, you are still tempted to guess at results, especially if the alternative is yet another marathon session with a pocket calculator. With TCALC you can *know* the answer in a couple of minutes—and find out (as I did) how poor your guesses usually are!

TCALC and the ideas behind it are all analytical techniques, and you have to create an initial design before you can start to analyse and optimize it. It is a lot easier to optimize a good initial design than a bad one, and Part 4 will show that the main difference lies in good or bad gain distribution.

Notes and references

[1] The assumption that ips will add in phase (coherently) is pessimistic, so any surprises should be pleasant ones. Coherent addition will apply either in wideband systems, where incidental phase-shifts are low, or (as here) where the frequency separations are small.

[2] *Introduction to Radio Frequency Design,* W A Hayward (Prentice-Hall Inc, Englewood Cliffs, New Jersey, 1982, ISBN 0-13-494021-0), Chapters 1 and 6.
NB. There is a misprint in the formula for the current-limited output; the formula given here is correct and so is the rest of Hayward's text on the subject.

[3] A significant exception where voltage limiting may apply is the case of the dual-gate GaAsfet. Unlike silicon dual-gate mosfets, or single-gate GaAsfets, dual-gate GaAsfets have Schottky gates which will conduct if forward biased. John Regnault, G4SWX, has pointed out that owing to the physical layout of the gates on the chip, forward conduction of Gate 2 can limit the peak signal at the drain to 2V or less. In this case, terminating the drain in a high impedance will make voltage limiting *worse* rather than better.

[4] This generalization has been immortalized by G4DGU as "Bartram's Law". (*TT, Rad Com* April 1984, p315). Witness a recent description (in another magazine) of a dual-gate GaAsfet preamp with a "third-order intercept" of +20dBm. You shouldn't need either me or G4DGU to tell you that this intercept was referred to the *output* of the preamp: subtracting the gain of 24dB left the input intercept a miserable −4dBm!

[5] Peter Chadwick, G3RZP, has surveyed a large number of ring-mixer data sheets, and found that the third-order input intercept is generally about 5dB above the lo level, and the 1dB gain compression point is about 5dB below it. However, achieving the potential intermodulation performance of most diode-ring mixers depends very much on proper termination of the input, lo and i.f. ports.

[6] "Modern Receiver Mixers for High Dynamic Range", Doug DeMaw, W1FB, and George Collins, AD0W. *QST* January 1981, p19. Hayward gives a good description of diode-ring switching mixers in Chapter 6 of *Introduction to Radio Frequency Design.*

[7] "Dynamic Range, Intermodulation and Phase Noise", Peter Chadwick, G3RZP. *Rad Com* March 1984, p221.

[8] Normally it is convenient for TCALC to assume automatically that passive attenuators (CAB/ATT stages) contribute no intermodulation. Intermodulation in such stages can still be introduced by making TCALC treat them as AMPlifiers. For example, a crystal filter with an insertion loss of 3dB can be described as an AMP having a noise figure of 3dB and a gain of −3dB. The result of the noise analysis remains exactly correct, the only penalty being that the filter is given the inappropriate label "AMP", but it now becomes possible to input an intermodulation intercept.

TO BE CONCLUDED

MODERN VHF/UHF FRONT-END DESIGN

Ian White, G3SEK*

PART 4. PRACTICAL FRONT-END SYSTEMS

IN THIS FINAL PART I will look at some typical amateur vhf/uhf receiving systems. This time I won't burden you with any maths at all. I've already given you those tools in Parts 1-3, including the TCALC program to take the drudgery out of the calculations. I've also explained how to feed sensible assumptions into TCALC so that you can trust the results. Some of my overall conclusions about typical amateur receiving systems may not go down too well in certain quarters, but they're all based on data from published equipment reviews and backed up by calculations which you can check for yourself.

I have already said that most commercial transceiver front-ends are noticeably deaf—as any vhf/uhf dxer knows! Reviewers have found that the noise figures of most 144MHz transceivers run in the range 3-7dB, though they can vary widely among different samples of the same rig, depending perhaps on whether they were aligned on a Friday afternoon . . . Since the target nf I recommended in Part 1 was about 2·2dB (noise temperature = 200K), this kind of performance leaves a lot to be desired. What can we do about it?

How to choose a preamp

One of the most common things people do to improve their deaf transceiver front-ends is to add a preamplifier. The first thing to do when choosing a preamp is to decide why you want it. Here are some possible reasons:

1. Your existing transceiver is deaf—the most common reason.
2. You want to be able to winkle out the weakest dx while the band is quiet and there aren't any strong signals about.
3. You have large, permanent, unavoidable feedline losses.
4. You work Oscar 10 or moonbounce, so you can take advantage of the lower antenna noise temperatures when pointing skywards.

In the last two cases you can generally get away with a permanently-connected preamp. However, in the first two cases, which are far more common, the preamp will make your system vulnerable to overload by strong signals [1]. Earlier parts of this article explained the different overload mechanisms in more detail. What generally happens is that the gain of the preamp provokes overload in the existing transceiver front-end; it is quite rare for the preamp itself to overload.

There are two cures for overload provoked by a preamp, and you can use both . The first is to use a switchable preamp, so you can get rid of it when there are strong signals about. With a bit more effort you will probably still be able to copy all but the weakest dx, which is far better than having the whole band wiped out by receiver overload! The second cure is to use the least possible amount of preamp gain, and I'll come back to that later.

Having decided whether the preamp needs to be switchable (and it probably does) the next step is to look at some advertised preamp specifications, particularly at noise figure and gain. When you add a preamp to an existing transceiver, the new system nf obviously depends to some extent to on how good or bad the transceiver was in the first place. Since the preamp manufacturers don't know the nf of your particular transceiver, they have no choice but to quote the nf of the preamp alone, as if it were connected to an idealized noiseless receiver. So the quoted nf is always better than you can get in practice. Nobody's trying to deceive you; you simply need to be aware of what the advertised specification really means.

The next thing to look at in preamp specifications is the gain. I'll say it once more: *too much gain is bad.* A typical 144MHz preamp with an nf of 1·3dB and gain of 15dB, as advertised, will improve the 7dB nf of a typical transceiver to about 1·7dB (Fig 17). Fair enough, it will transform a deaf receiving system into one which, allowing for feedline losses, comes pretty close to our target nf of around 2dB. But that 15dB of preamp gain will also transform the strong-signal performance from mediocre to downright bad. That is the penalty of improving the system nf by adding things on to the transceiver rather than tackling the real problem, the existing poor front-end.

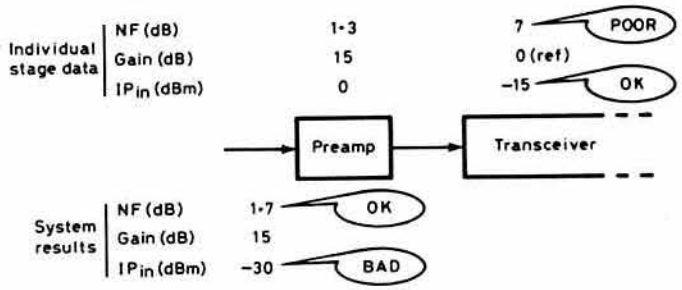

Fig 17. Effect of a preamplifier on a typical commercial 144MHz transceiver

Reputable preamp manufacturers are actually somewhat apologetic about the fact that their products need to have some gain! Manufacturers who boast about the gain of their preamps either have little regard for your technical knowledge, or they don't understand what they're doing [2]. Unfortunately there are always some customers who won't believe a receiver is any good unless the S-meter is perpetually bouncing about and the speaker pours out noise like Niagara Falls! These are the same people who believe in the magic powers of this years's fashionable devices, expecially if they're GaAsfets . . . So read the preamp specifications carefully; and then go over them again, this time reading *between* the lines!

Since the preamp itself is unlikely to overload, except in rather unusual situations, you don't need to pay too much attention to its stong-signal performance. Even so, it's better to buy from a manufacturer who gives some evidence of having considered the problem in the first place, and can give you some support if you do happen to live close by a tv/fm transmitter [3]. Remember these points about intermod performance from Part 3:

(a) *Input* intermodulation intercept is what really matters in a system.

(b) Output intercept looks more impressive, but you have to subtract the gain to get the input intercept.

(c) Input intercepts significantly above 0dBm are unlikely unless the preamp uses special techniques such as negative feedback.

And finally:

(d) Don't be besotted with low noise figures; don't be swayed by this year's fashion; and don't believe in magic!

Adjusting preamp gain

Having fitted a preamp, you can take further precautions against overload by setting the gain to the minimum that is necessary to get your system nf down. Vastly excessive gains (20dB and more) signify bad design, but even well-designed commercial preamps tend to have plenty of gain in order to produce some improvement on even the deafest transceiver. Your transceiver may not be that bad, so you can get rid of the excess preamp gain. Some preamps already include a variable attenuator at the output, so that's a good start [4]. If yours does not, you can easily fit one following the circuit of Fig 18. Take care if the preamp includes rf switching or has a dc power feed up the coaxial cable; if in doubt, check with the technical guru at your local radio club.

It's very simple to adjust the preamp gain down to the minimum your system needs. Start with full gain and tune in a weak and slightly noisy fm signal. Reduce the preamp gain until the audio starts to become noisier, and then increase the gain *very slightly* until you no longer notice the deterioration [5]. That's all it takes: if your system still gets crunched by strong signals, at least you know you've done your best!

What's wrong with commercial transceiver front-ends?

The front-ends of transceivers from the "Big Three" manufacturers are not sensitive enough for vhf or uhf dx working, and they are none too happy with strong signals either. We know this from equipment reviews and from personal experience. If you fit a preamplifier to improve the sensitivity, the strong-signal performance becomes even worse. Yet we also know that it *is* possible to build front-ends which outperform those of commercial transceivers on both weak and strong signals. So, why can't we buy these front-ends from the "Big Three" manufacturers? In their hf models they have recognized our requirements for strong-signal performance as well as sensitivity, and "high dynamic range" is now a strong selling point. But at vhf and uhf all the recent design effort seems to have gone into digital features like scanning, memories, voice synthesizers and now (Lord help us)

* 52 Abingdon Road, Drayton, Abingdon, Oxon OX14 4HP

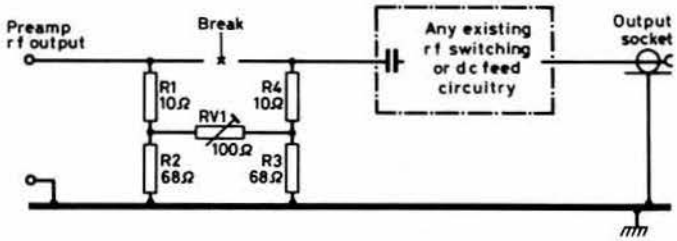

Fig 18. Fitting a 5–15dB variable attenuator to the output of a vhf preamplifier. RV1 is a miniature cermet or carbon trimpot, and all other resistors are miniature carbon film or metal oxide. *NB:* If the preamp has rf bypass switching or a dc feed via the coaxial output connector, make sure that the attenuator never carries rf or dc power

Ascii code identifiers. These are all very well in their way—but what about the radio bit?

Front-end design in commercial vhf and uhf transceivers is still stuck with the dual gate mosfet rf amplifier and mixer—technology over a decade old, and all too reminiscent of the previous decade's valve designs. That was the situation back in 1977 when I made the very first replacement front-end board for my own FT221, and it's hardly any better today. Front-end performance of commercial vhf/uhf transceivers has merely crept forward instead of taking the big strides it should have. New and potentially useful devices have been either ignored or, like GaAsfets, wasted by being plugged into the same tired old circuitry.

Without significantly increasing the price of a top-of-the-line transceiver, it should be perfectly possible for the "Big Three" to mass-produce modern vhf and uhf front-ends with far better performance on both weak and strong signals. After all, they've already done it with their hf rigs.

At present I really can't recommend *any* commercial transceiver for serious vhf/uhf dx and contest working; not without a lot of modification. You can judge for yourself about general operating facilities and the "feel" of the rig. My recommendations here are based mainly on how easy it is to completely replace the front-end, either by exchanging a circuit board or by finding space for a new front-end somewhere inside the box. On those grounds I recommend the IC271, the IC251 and the FT225 as fairly recent transceivers that can be turned into serious 144MHz dx rigs by replacing the existing front-ends. Thus modified, these rigs set the standard by which others can be judged [6]. The older FT221 and IC211 have somewhat higher local oscillator (lo) noise and fewer general facilities, but after modification even these are way ahead of the newest off-the-shelf transceivers when it comes to dynamic range. And they'll stay ahead until the "Big Three" get their acts together.

Transverters revisited

VHF and uhf ssb started in the 'sixties when hf ssb enthusiasts added transverters to their hf transceivers. This opened up the vhf bands to the advantages of transceive operation, vfo control, stable and selective receivers, and of course to the advantages of ssb as a dx mode. But by the late 'seventies transverters had acquired a bad reputation for receiver overload, and on 144MHz they have been almost entirely superseded by single-band, single-conversion transceivers for serious dx and contest work. However, transverters have always been widely used on the other vhf/uhf/microwave bands, and now they are making a comeback on 144MHz. So what went wrong with transverters in the 'seventies and what's changed since then?

A transverter followed by an hf transceiver is a multiple-conversion system. The receive signal path in the transverter is very like the front-end of a normal single-conversion transceiver, except for one thing. The lo in a transverter is crystal-controlled on a fixed frequency, and the hf transceiver does the tuning. That means you can't put a narrowband i.f. filter directly after the transverter front-end, so the hf transceiver front-end is vulnerable to overload from the amplified wideband output of the transverter.

Transverters are worth a second look today because the front-ends of commercial hf transceivers have improved in the past few years, while those of vhf/uhf transceivers have not. Also the i.f. performance of hf transceivers has always been a few steps ahead; even today, few vhf/uhf transceivers can offer adjustable noise blankers, switchable filters with really good stopband rejection, variable bandwidth, i.f. shift or notch filtering, all of which could be handy on today's crowded vhf and uhf bands.

The techniques for analysing the performance of a single-conversion front-end will work equally well for a transverter system. Let's start by seeing what went wrong in the 'seventies. At that time, transistorized hf transceivers were just coming in, and we all thought they were wonderful.

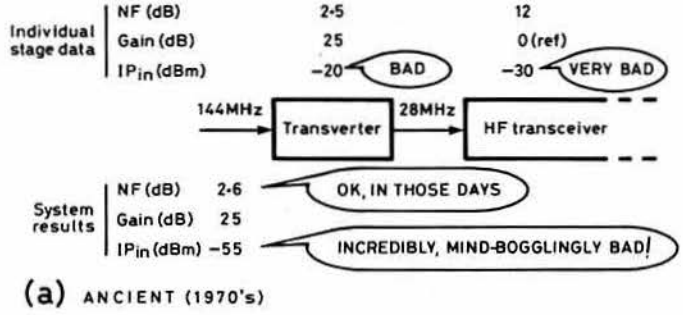

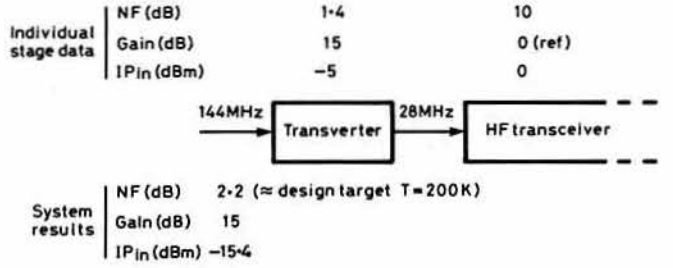

Fig 19. 144MHz transverters with hf transceivers

Actually, their strong-signal handling was rotten: their third-order intermodulation intercepts were −30dB or worse [7]. A 'seventies-style transverter (which you can still buy today!) had about 25dB of gain, which dropped the system input intercept at 144MHz to −55dB or lower. With a system noise floor of −140dBm the dymanic range would be less than 60dB (Fig 19(a)). In practical terms, signals over about S9 anywhere on the band would cause havoc!

Now do the same analysis for a more modern system. A typical nf for an hf transceiver at 28MHz would be 10dB, with a third-order input intercept of 0dBm. Using a transverter designed along the same lines as the front-end in Fig 5 (Part 1), the gain required to achieve a system noise temperature target of 200K would be about 15dB. Assuming there is no intermodulation in the transverter, the system input intercept would be about −15dBm (Fig 19(b)). That's not as good as a single-conversion front-end designed with equal care, but it's a lot better than an ordinary 144MHz transceiver plus preamp. And unlike a single-band 144MHz transceiver, the hf-plus-transverter approach equips you for the hf bands and allows you to add further transverters for 50MHz, 70MHz and 432MHz, or even 1·3GHz. Worth thinking about?

If so, you'd better think about these points too. First of all, you can only get good vhf/uhf strong-signal handling if the hf transceiver itself is good in this department. It's well worth reading the reviews and shopping around for a model which combines a noise figure of about 10dB [8] at 28MHz with a third-order input intercept of 0dBm or better. Second, if you do have an hf transceiver that good, it will show up the deficiencies in transverters which still use the old dual-gate mosfet front-end technology and may have input intercepts of −20dBm or worse. We have already seen that a decent 144MHz front-end with a diode ring mixer can have an input intercept of −5dBm or better. With robust rf and i.f. amplifier devices and high level ring mixers *in the right system design,* the transverter input intercept should definitely be better than 0dBm. A third point to remember is that you can use an attenuator just like the one in Fig 18 at the 28MHz output of the transverter to keep down the gain. The gain of an off-the-shelf transverter will be enough for the deafest hf transceiver, and will be more than you really need if you have chosen your own hf transceiver well.

Turning briefly to the bands other than 144MHz, a target noise temperature of 150K at 432MHz and above could be achieved using a transverter receive gain of no more than 15–20dB. To keep the overall gain this low, a masthead preamplifier would be essential on these higher bands. By the way, don't imagine that there are no problems of front-end overload on the microwave bands; things may be quiet normally, but just wait until a strong portable station comes along and takes your front-end apart for an entire weekend! The "pile on the preamps" attitude is still very strong on the microwave bands, and many preamp/converter systems would benefit from a going-over using TCALC to see how much front-end gain you could manage without. For example, a modern dx front-end for

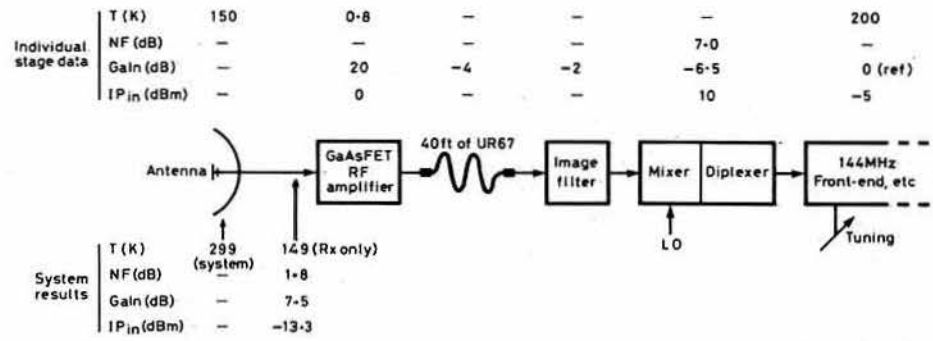

Fig 20. Suggested 1·3GHz dx front-end

1·3GHz could consist of a GaAsfet masthead preamp, a run of UR67 cable down to an image filter, and a good diode mixer feeding a 144MHz transceiver—one with a modern front-end, of course (Fig 20)! Here, just for once, you do need a GaAsfet running at full gain; but you don't need any other rf amplifier, and the overall system gain is still low.

Antenna noise temperatures on 70 and 50MHz are much higher than on 144MHz or above, so the requirements for low receiver noise temperatures are considerably relaxed and the transverter receive gain can be reduced accordingly. On these bands you should be seeing transverter input intercepts better than 0dBm.

Notes and references

[1] A few months ago I was mildly amused to see someone stating in another magazine that although a preamp *theoretically* degrades strong-signal performance he'd noticed no actual signs of it. If he worked dx on the same vhf and uhf bands as you and I, he'd notice it all right!

[2] Reputable preamp suppliers advertise performance figures which are routinely achieved in manufacture, and they can back their claims by measurement. This applies particularly to noise figures, which are not at all easy to measure properly. Less reputable firms seem to advertise the best nf they ever measured, or hoped for; or worst of all merely quote the device nf from the transistor manufacturer's datasheet. The last of these is particularly easy to spot, and generally signifies that the preamp manufacturer hasn't much of a clue, or much test equipment either. Read "The effects of preamplifiers on receiver performance, and a review of some currently-available 144MHz preamplifiers," by Julian Gannaway, G3YGF (*Rad Com* November, December 1981, pp1026, 1120) and draw your own conclusions.

[3] I mentioned in Note [1] of Part 2, the situation where second-order intermodulation products of tv/fm broadcast transmissions fall into the amateur band. If the problem only appears when you add a preamp to your existing transceiver, it's almost certainly due to your preamp having appreciable gain on the broadcast frequencies outside the amateur band, causing second-order intermod in the transceiver itself. It's not likely that the intermod is occurring in the preamp itself. The normal cure is to put a filter between the preamp and the transceiver, or to use a preamp which has a filter built-in.

[4] Preamps with adjustable *input* attenuators are not satisfactory, because attenuation at the input adds directly to the nf. Also avoid preamps which vary the gain by biasing off the transistor; this is a sure invitation to overload in the preamp itself. Look for a preamp with a resistive attenuator at the output, similar to Fig 18.

[5] By making the adjustment on a noisy fm signal you are taking advantage of the "threshold effect" of fm detection. If a signal is weak enough to sound noisy, then even a small decrease in rf signal:noise ratio will make the audio s:n ratio very much worse. This makes a small change in rf s:n ratio quite easy to detect by ear without the need for instruments. (Thanks to muTek Ltd for this tip.)

[6] Peter Hart, G3SJX, prefers the FT225 with muTek front-end (*Rad Com* April 1984, p312) while Angus McKenzie, G3OSS (*Amateur Radio* September 1984) recommends the more recent IC271E, again with a muTek front-end. Other transceivers, even different models from the same manufacturers, are not so convenient to modify.

[7] The third-order input intercepts of various unmodified early models of the FT101 (before the "Z" models, which are quite different internally) have been quoted in the range −30dBm to worse than −50dBm! Many other transistorized hf transceivers of that vintage have equally fragile front-ends.

[8] The sensitivities of hf transceivers are usually specified in terms of microvolts for a given s:n ratio, rather than noise figure. Note [5] of Part 1 explains how to convert between the two systems. A noise figure of 10dB corresponds to a noise floor of about −130dBm (in a 50Ω system with 2·5kHz noise bandwidth); or 0·20μV pd for 10dB s+n:n (as quoted in RSGB reviews); or 0·26μV pd for 12dB sinad (as quoted in reviews by G3OSS). For transceivers with switchable hf preamps, take care to compare sensitivities and input intercepts under the same conditions. Also note that intermodulation performance on 28MHz is often not as good as on the lower bands.

SUMMARY OF THE SERIES

1. A receiver front-end can be designed as a complete system to deliver the optimum overall performance. When you have the system design right, you can then look for circuits to go in the boxes in the block diagram.

2. For ordinary terrestrial dx and contest working, recommended targets for receiver noise temperature are 200K (noise figure around 2·2dB) for 144MHz, and 150K (nf around 1·8dB) for 432MHz and above. These are roughly equal to antenna noise temperatures on the respective bands.

3. Front-ends with lower noise temperatures than recommended above will be unnecessarily vulnerable to overload by strong signals; you may end up hearing *less* dx, not more.

4. There are three main types of front-end overload: gain compression, intermodulation (especially third-order) and reciprocal mixing. The first two are problems in the signal path; reciprocal mixing is caused by noise from the local oscillator. All these overload mechanisms can cause interference to weak wanted signals.

5. Excessive front-end gain is the biggest enemy of good strong-signal performance. Front-end gain should be no higher than is necessary to achieve the target noise temperature. Low-noise devices are useful because they allow you to meet your target for receiver noise temperature while using less overall front-end gain: that is the correct role of low-noise devices in modern front-ends.

6. You can predict the sensitivity and intermodulation performance of front-ends using a home computer program like TCALC. I have shown how to feed sensible assumptions into TCALC so that you can trust the results. By using TCALC you can experiment with the system design and get it right before you even begin to think about detailed circuit design. Even if you buy everything ready-made, TCALC can help you assemble the best possible overall system.

7. With the right system design, quite ordinary components and devices will give good results. With the wrong system design, even magic words like "GaAsfet" won't save your front-end from mediocrity. Good system design is particularly important if you don't have access to advanced test equipment.

8. It is possible to design and build vhf and uhf front-ends which meet the above targets for noise temperature, and also are virtually immune to overload effects—even at "open-contest" signal levels. You are entitled to expect this kind of performance from a vhf/uhf receiver front-end.

9. The front-ends of vhf/uhf transceivers from the "Big Three" manufacturers do not measure up to these standards. I cannot recommend *any* commercial transceiver for serious vhf/uhf dx working unless its front-end is replaced.

10. Preamplifiers are a second-best approach to reviving deaf transceivers. If you choose and use them wisely they can give acceptable results for casual operating.

11. A good hf transceiver plus a modern high-performance transverter is well worth considering as an alternative to a single-band transceiver.

12. To get the best out of a system of ready-made units, you must buy wisely and be prepared to make adjustments. It is always worthwhile to adjust the gain of a preamplifier or transverter down to the minimum that your system actually requires.

Acknowledgements

Thanks to G3OSS, G3SJX, G3YGF and the ARRL for equipment reviews which are a valuable source of data; to W1FB and W7ZOI for their books and articles; and to G3RZP, G4PMK, G4SWX and especially G4DGU for all their frank and friendly comments. □